Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles
Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Cover Story

The Next 10 Years: a Challenge for the Industry

Louis J. Burns
Vice President, General Manager
Desktop Platforms Group
Intel Corporation

Overview

Our industry has come a very long way in the past 20 years.

The starting point was the original IBM PC*, with its 8-bit CPU, kilobytes of main memory, simple floppy disk storage, and first-generation bus and I/O architecture. It was a great machine for running VisiCalc*, the “killer app” of the day.

Today’s advanced PC platforms are based on the Pentium® 4 processor, Intel’s highest performance processor for desktop PCs. These systems are at the center of the digital world. They provide performance where today’s users need it most, and are designed for where the Internet is going.

The next 10 years will see advances in PC performance, applications development, and user experiences at least as dramatic as what we’ve experienced since 1981, including the ability to deliver rich media on demand to any user, anywhere, at any time.

The industry faces two immediate challenges. The first is how to make wireless connectivity pervasive on all PC platforms. The second challenge is the creation of a third-generation I/O architecture whose bandwidth can cost-effectively scale with coming advances in processor performance.

Overarching these technical challenges is the requirement to maintain the standards-based architectural paradigm that is essential for the PC to remain the volume platform. Intel is now at work within the industry to help meet each of these challenges.

Volume Platform

The PC has undergone a tremendous evolution in the past 20 years. The true power of the PC platform is that it has established a homogeneous computing environment all with similar hardware capabilities, a standard operating environment and an installed based measured in hundreds of millions of units.

As a result it created a scalable, cost-effective platform architecture that soon became pervasive throughout the PC industry.

This homogeneous environment set the stage for continuous waves of application development that helped drive the evolution of the PC with capabilities we see emerging today: Pentium 4 processors capable of 2-GHz, the 1-Gbyte/sec. system bus, the 1.6-Gbyte/sec. memory bus, and 1-Gbyte/sec. I/O.

At Intel we think the best is yet to come. In the next few years, wireless technologies will become pervasive in PC platforms, both desktop and mobile, enabling PC connectivity to scale across multiple PC, communications, and consumer electronics market segments. In addition, PC platform technologies will continue to evolve to encompass 10-GHz processors, and like improvements in memory and I/O technologies.

Pervasive Wireless

The arrival of the "Extended PC Era," including the convergence of personal area network (PAN), local area network (LAN), and wide area network (WAN) wireless technologies, is creating unprecedented new market opportunities for the industry.

It is clear that to make pervasive wireless connectivity a reality, no single radio technology will fit all the requirements of every platform and user. We will need multiple ways to connect in order to achieve the optimum combination of bandwidth, reliability, and cost for different usage scenarios and applications (see Figure 1).

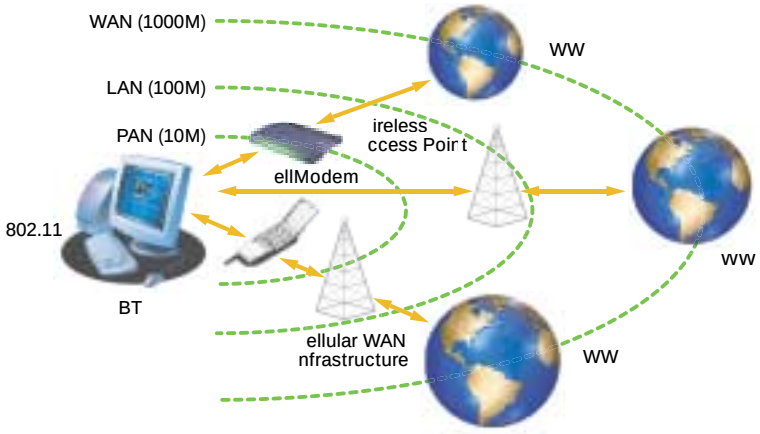

Figure 1. Intel® Wireless Vision: Pervasive Wireless Connectivity

Examples include:

- Bluetooth* wireless connectivity for personal area network (PAN) out to 10 meters.
- Wireless Ethernet (IEEE 802.11) enabling wireless LANs out to 100 meters.
- Cellular technologies in the wireless WAN infrastructure and wireless access points providing connections to the Web.

Intel plans to deliver native wireless technology in our platforms within the next few years. To make this a reality, Intel is currently working with the industry, standards bodies and regulatory authorities to help overcome the challenges of multiple standards, interoperability, security, regulatory requirements, and cost.

I/O Architecture Evolution

To continue to evolve the PC platform Intel has identified the need for a third-generation I/O architecture (see Figure 2), successor to ISA and PCI. Like past efforts within the industry to advance platform initiatives such as PCI, AGP, Intel® Hub Architecture, and USB for external I/O, Intel will work with the industry to deliver an open specification that addresses these challenges. Typically, I/O architectures are a big investment for the industry thus it is critical that a new direction looks over a 10-year horizon and meets demanding requirements for the long term. The goals of this effort must address the challenges of flexibility, connectivity, scalability, and cost that have been crucial to the PC platforms evolution.

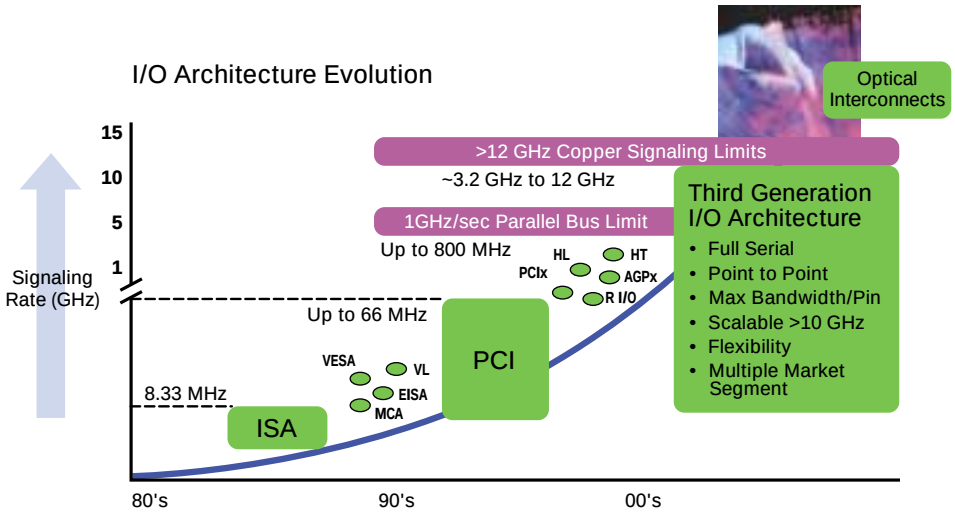

Figure 2. I/O Architecture Evolution

While the PCI bus has been notably successful, it is becoming clear that multi-drop parallel bus technology such as PCI is reaching its limits. With PCI, 70 pins are required to deliver a bandwidth of slightly more than half a Gbyte/sec.

A variety of new interconnect solutions have been put forward but what is needed is a true serial bus technology for scalability well into the future, up to the limits of copper signaling technology.

As bandwidth demands continue to grow, the industry clearly needs a new interconnect technology capable of delivering much greater bandwidth per pin efficiency. The goal of this new I/O architecture is to enable system designers to scale bandwidth by adding serial ports. It will be cost-effective for desktop PCs, power-friendly for mobile systems, and deliver high bandwidth needed for demanding applications such as communications.

This open specification is under development now. A preliminary version will be available for your review in fall 2001.

Summary

Dramatic gains in processor performance and Internet bandwidth will certainly continue in the next 10 years. Can PC I/O technologies keep pace? More importantly, can PC interconnect technologies scale while maintaining the homogeneity and the open standards environment needed to keep the PC a high-volume platform?

The challenge for the industry will be to develop a third-generation I/O architecture that is scalable across industry market segments and capable of taking bus performance to the limits of copper technology, and then beyond, to optical performance levels.

Author Bio

As vice president, general manager Desktop Platforms Group, Louis J. Burns is responsible for design, development, and bringing to market Intel's desktop platform solutions, including processors, chipsets, motherboards, software and services. He shares these responsibilities with William Siu, vice president and general manager, Desktop Platforms Group.

Most recently, Louis was vice president and general manager, Platform Components Group, the primary producer of core logic and integrated graphics chipsets for Intel. He also served as vice president and director of Intel's Information Technology group for four years, supplying computing capability to Intel worldwide. In this role, he learned what IT shops struggle with day-to-day, from strategic decisions for the company to tactical product deployment challenges.

Previously, Louis spent 12 years in Intel's Sales, Marketing and Applications groups, gaining the sales perspective of the organization. From seller to customer, Louis has the breadth and depth of experience for a rapidly evolving global marketplace. He was appointed an Intel vice president in 1996 and elected a corporate vice president in 1997.

Column

From the Editor
Donna Loveland
Managing Editor
Intel Developer Update Magazine
Intel Corporation

Column

At the Intel Developer Forum Conference, Spring 2001 the computing industry explored the theme "Expanding the Power of the Net." This month's issue, fresh from the February 26–March 1 event, includes a selection of the topics explored in more than 250 labs and sessions.

The Next 10 Years: a Challenge for the Industry—cover story—To capture new opportunities for growth during the next 10 years, the PC industry faces the twin challenges of making wireless connectivity pervasive on all PC platforms, and creating a scalable third-generation interconnect bus architecture.

Building Blocks of Wireless Internet Clients—The Intel® Personal Internet Client Architecture helps accelerate development of hardware and software solutions by partitioning the client platform into memory, communications, and computing subsystems.

Making Wireless Networking Easier For Consumers—Wireless home networking products are providing reasonable levels of performance and functionality, but usability lags behind. And usability is crucial in driving consumer acceptance.

InfiniBand* Architecture-Enabling Technology—InfiniBand* architecture enabling technology is available now, including product development kits, port logic, and target transport services.

Innovative PCs Deliver Performance, Simplicity, Personality—Thanks to widespread industry innovations, PCs are now capable of delivering unprecedented performance from a convenient and attractive package.

Using SSE and SSE2: Misconceptions and Reality—Programming with SSE/SSE2 is quick and easy. It uses mature tools and greatly improves performance on a wide variety of multimedia and communications applications.

Intel Enhances the Intel® IXA™ Processor Family—Intel recently announced major enhancements to the IXA™ network building blocks: a new, faster network processor, an enhanced SDK that includes a library of 100 macros, and the Intel® Microengine C Compiler.

Serial ATA—the Long-term Solution for Storage Connectivity—Serial ATA delivers the best long-term storage interface solution, addressing the shortcomings of parallel ATA while delivering a scalable interface solution supporting several speed doublings.

Cross-Platform Remote Server Management Capabilities—Version 1.5 of the Intelligent Platform Management Interface (IPMI) specification enables development of cross-platform software for remotely managing servers and other mission-critical systems.

Intel® Memory Technology Update—Memory technologies and market dynamics are examined in the context of the Intel® Desktop Memory Roadmap.

Volume II of Intel® Developer Solutions Catalog Now Available—Volume II of the *Intel Developer Solutions Catalog* is a concise reference resource to help software developers quickly find tools and applications optimized for the Intel® Architecture.

Intel® Motherboard Selector Guide Goes Online—Redesigned to include more motherboard vendors and products, more granular search features, and an easier interface, the Guide is a simple way to find information on all Intel® Architecture-based motherboards on the market.

SoundMAX* 2.0 Delivers Low-Cost State-of-the-Art Sound—The SoundMAX 2.0 is an integrated audio solution that enables integrators to add the highest quality sound capabilities to Intel motherboards at a very low cost.

Next month, be sure to visit *Intel Developer Update* for continuing in-depth information from the IDF Spring Conference.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Technology and Initiatives Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held communications positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in expository writing from the University of Iowa.

Departments

Desktop

Intel® Motherboard Selector Guide Goes Online

Joseph Yu
Product Marketing Engineer
Desktop Platform Group
Intel Corporation

Overview

The Intel® Motherboard Selector Guide is a comprehensive, searchable, Web-based compilation of all currently available Intel® Architecture-based motherboards. It is an extremely powerful yet user-friendly Microsoft SQL* 7 database encompassing more than 700 motherboards and more than 50 motherboard vendors. The Selector Guide has been redesigned with more granular search features and an easier interface to help integrators quickly find motherboards according to the criteria that matters to them. The Motherboard Selector Guide is actively supported and maintained so that the latest information is in the database. Currently, it is the only Web-based compilation of its breadth and kind.

Selector Guide Cuts Down on Motherboard Research Time

With the numbers of motherboards on the market with different CPUs, chipsets, integrated features, and form factors, system integrators and OEMs sometimes face the dilemma of having too many to choose from. It can take three or four hours to search on the Internet and find motherboards with the features they need, such as compatibility with the Intel® Pentium® 4 processor. It can take even longer to browse through trade catalogs. Until now there has been no easy way to find all motherboards that fit a certain criteria. Using the Motherboard Selector Guide, the search time is narrowed to seconds.

Originally launched at the beginning of 2000, the Motherboard Selector Guide has been re-launched with more granular search features, an easier to use interface, and an expanded pool of motherboards and vendors. A user can do a basic search to identify a broad range of motherboards from Intel and 49 other motherboard vendors, including the industry leaders. Or the user can find a specific motherboard by providing detailed criteria in the advanced search method.

The Motherboard Selector Guide offers instant knowledge and ability to compare the products in the market place. For example, suppose a system integrator or home user wants to utilize an external graphics solution but does not want to give up all the innovative features and proven reliability of the Intel® 815 chipset family. With a click of the mouse, Motherboard Selector Guide lists most of the motherboard solutions in the market today that are equipped with Intel® 815EP chipset.

What if a PC enthusiast wants to build a system with a P4 processor to speed up his or her MP3 or 3D graphic decoding but does not want to throw away that old ATX chassis? Again, the Motherboard Selector Guide offers choices to limit the search result based on the processor type (Pentium 4) and the form factor (ATX).

Expanded Search Features Enable Broad or Narrow Searches

The Intel Motherboard Selector Guide has a comprehensive set of search features that let integrators research motherboards as broadly or as narrowly as they need to. They can select motherboards based on the following criteria:

- *Intel® processor* (Intel Pentium 4 processor, Intel® Pentium® III Xeon™ processor, etc.)
- *Intel® chipset* (Intel® 850 chipset, Intel® 820E chipset, etc.)
- *Form factor* (ATX, ExtendedATX, MicroATX, etc.)
- *Manufacturer* (more than 50 leading vendors)
- *Integrated features* (graphics, audio, LAN, SCSI, etc.)
- *Processor package* (PGA 423, Slot 1, Socket 370, etc.)

An example of a search page from the Selector Guide appears in Figure 1.

Figure 1. Intel's Motherboard Selector Guide has comprehensive search features and an easy reference.

The power to search does not stop here. Visitors are able to find even more detailed information by clicking the Web site links available in the search results. For example, if someone is interested in utilizing a six-channel audio Communication and Networking Riser (CNR) Card and would like to know which Pentium 4 motherboards support the CNR, the Selector Guide can provide linkage to the technical data maintained by specific vendors and satisfy the inquiry.

The Motherboard Selector Guide Team welcome any feedback from the users, the database is very flexible and is scalable to include additional features (such as BIOS version and compatible memory type) that might be important to the end users. Please e-mail us with your feedback.

Commitment to Database Maintenance

The 50 vendors listed in the Motherboard Selector Guide represent more than 90 percent of motherboard manufacturers. Although the Selector Guide is owned and hosted by Intel, it is capable to be maintained directly by the motherboard vendors. As new motherboards come into production, the vendors will be able to update the database in real time.

The motherboard vendors' administration section shown in Figure 2 provides a simple and efficient interface for motherboard manufacturers to maintain and update their product information. It takes less than a minute to create a new board entry in the Selector Guide and no more than 30 minutes per week for each vendor to maintain its product database.

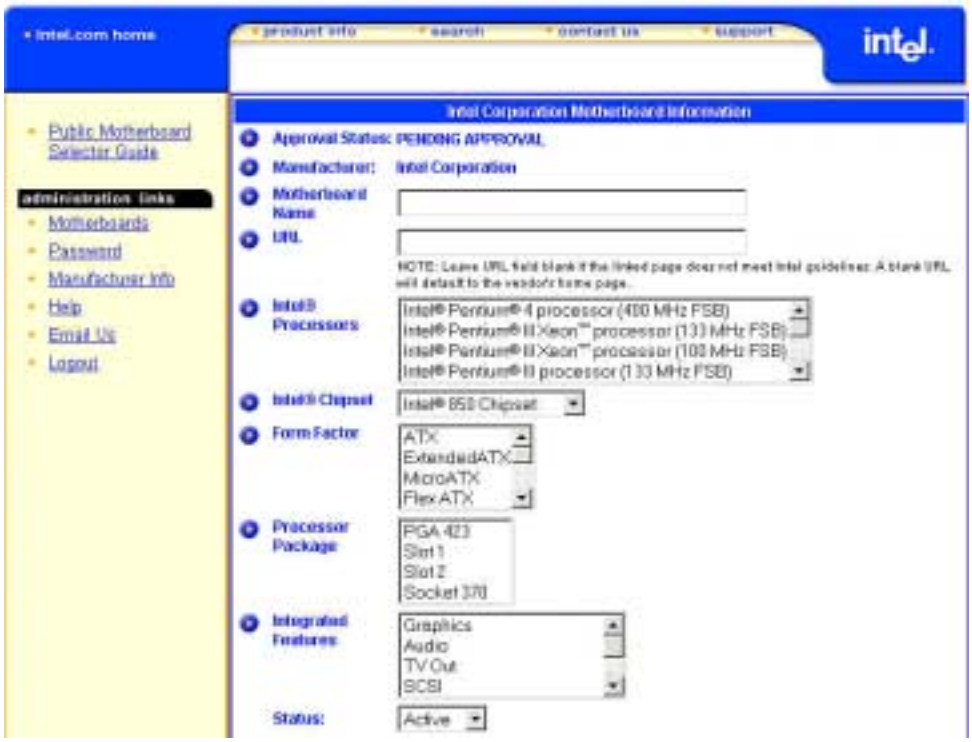

Figure 2. The Intel® Motherboard Selector Guide provides vendors a friendly and efficient administrator interface.

Summary

The Intel Motherboard Selector Guide is an extremely powerful yet user-friendly Web-based compilation of all currently available Intel Architecture-based motherboards currently produced by leading vendors. It enables an extremely refined search according to very narrow criteria, reducing research time to minutes. It also provides an opportunity for system integrators and home users to access a broader range of selections based on their specific need.

More Info

You can find the Intel Motherboard Selector Guide at the Intel Corporate Web site.

Author Bio

Joseph Yu joined Intel in 2000. Before joining the Desktop Platform Group, he worked in the Initiative Program Management Group as a platform application engineer. Joseph holds a B.S. in industrial management and an M.S. in public policy and management in information systems from Carnegie Mellon University.

SoundMAX* 2.0 Delivers Low-Cost State-of-the-Art Sound

Erik Cubbage
Product Marketing Engineer
Reseller Products Division
Intel Corporation

Overview

The use of sound on personal computers has become ubiquitous. Many users expect to use PCs to play games, to listen to CDs, or to experience MP3 streaming video and audio over the Internet. As a result, most PCs sold today incorporate audio capabilities. The SoundMAX* 2.0 AC'97 integrated audio solution offers some of the most advanced audio functionalities available today.

Integrated Audio Means Lower Costs

The SoundMAX 2.0 AC'97 integrated audio solution, developed by Analog Devices and available on Intel® Desktop Boards, includes high-performance codecs, application code, and drivers, delivers state-of-the-art audio capabilities. While add-in PCI sound cards offer similar capabilities, these cards typically require extra hardware and can cost up to $100 (USD). The SoundMAX 2.0 solution uses on-board components and the PC's processor to deliver high-performance sound. This enables integrators to add sound capabilities with fewer and less expensive parts, thereby reducing overall cost.

SoundMAX 2.0 is available for Intel® Pentium® 4-based D850GB™ as well as Pentium® III-based D815EEA™, D815EFV™, D815EPFV™, D815EEA2™, and D815EPEA2™ desktop processor boards.

Highly Compatible

To ensure users can employ the card with their favorite applications, the SoundMAX 2.0 is compatible with all popular Microsoft Windows* applications, including the DirectX* 8.0 multimedia codec. The device also supports the Musical Instrument Digital Interface (MIDI) standard, allowing users to simply plug and play their MIDI-compliant instruments.

Integrated Audio Controls

Included audio control software enables end users to adjust sound levels to account for the quality and position of their speakers. The simple and intuitive software-based controls are designed for ease of use, as well as for performance-minded consumers.

High-Quality Recording and Playback

The SoundMAX 2.0 solution offers a wide range of advanced recording and playback capabilities.

Variable sample rates. The SoundMAX 2.0 records speech at 8-, 11-, and 16-kHz sample rates. While the higher sample rates afford better quality, the lower rates require less CPU horsepower and disk space. Using these adjustments, users can select the most appropriate sample rate for their PC's capacity and desired audio quality.

Speech input. High-quality speech input features support applications that use voice control and support Internet telephony.

Rich, low-distortion playback. Low-distortion audio playback features furnish high-quality sound without hissing and popping. SoundMAX 2.0 plays recordings at a 20-kHz audio bandwidth that enables listeners to hear the full range of sound from the lowest base to the highest treble. This provides a richer sound than solutions that play back only a partial bandwidth.

Enhanced playback capabilities. The SoundMAX 2.0 offers additional features to further enhance the users' listening experience. The XG-Lite 360-sound sample set enables listeners to hear more instruments in an orchestra than they would otherwise. Channelized adjustable reverb and chorus provide special effects, such as echoes and reverberations. A CD-quality 44.1kHz output sampling rate allows users to hear individual sounds in a recording more distinctly.

Supports a wide range of file formats. Asynchronous record/playback supports any streaming or published audio file format with professional quality sound. The solution can also play back CD, DVD, MP3, or WAV files at any

independent sampling rate.

Environmental 3D Audio

Support for Sensaura 3D audio technologies means that users with two speakers or wearing a set of standard headphones will experience sound in three dimensions on most applications. The SoundMAX 2.0 currently supports the Digital Ear Head-Related Transfer Functions (HRTF) and the MacroFX* implementations of Sensaura 3D audio. It also supports ZoomFX*, which allows listeners to hear nearby sounds different than those far away. For example, ZoomFX would let you experience the sound of a bullet whizzing by your ear differently than a shot in the distance.

Easy to Upgrade

Because the SoundMAX 2.0 is primarily software-based, it enables scalability of the motherboard's audio subsystem for future upgrades. OEMs and integrators can upgrade to newer versions of SoundMAX via a simple software download without having to replace hardware on the PC's motherboard.

Summary

With the SoundMAX 2.0, users of Windows-compatible applications can take advantage of the most advanced professional-quality sound recording and playback capabilities to enhance their experience playing games, listening to CDs, or using Internet telephony. The SoundMAX 2.0's integrated audio technology means these high-end features are available at a fraction of the cost of PCI add-in solutions.

More Info

The SoundMAX drivers for the Intel® D850GB, D815EEA, D815EFV, D815EPFV, D815EEA2, D815EPEA2 boards are available for download at the Intel support site for consumer products.

For more information on the SoundMAX 2.0 product, visit the Analog Devices Web site.

Author Bio

Erik Cubbage joined Intel in 1996. He is currently a product marketing engineer in the Reseller Products Division (RPD) in Hillsboro, Oregon. Before joining the Reseller Products Division, Erik worked as a retail channel representative in Intel's field sales organization. Erik attended the University of Kansas.

Initiatives and Technologies

Intel® Memory Technology Update

Kuljit Bains
Memory Strategic Planning
Platform Memory Operations
Intel Corporation

Overview

Whatever the desktop market or price point, Intel is working with industry leaders to ensure developers enjoy a wide range of platform memory solutions. For developers seeking the highest system performance, for example, Intel offers high bandwidth Rambus* dynamic RAM (RDRAM) memory solutions that perfectly complement the Intel® Pentium® 4 processor and Intel® 850 chipset. For developers looking for respectable performance at a lower price, there is an array of cost-effective solutions using synchronous dynamic RAM (SDRAM) memory. In the next two years Intel will expand these options and add solutions that use the developing technology of double data rate (DDR) memory.

The Intel® Desktop Memory Roadmap

Intel is closely involved within the industry to make sure that robust platform solutions become available for every technology. The Intel® Desktop Memory Roadmap (see Figure 1) illustrates Intel's expectations for the next two years.

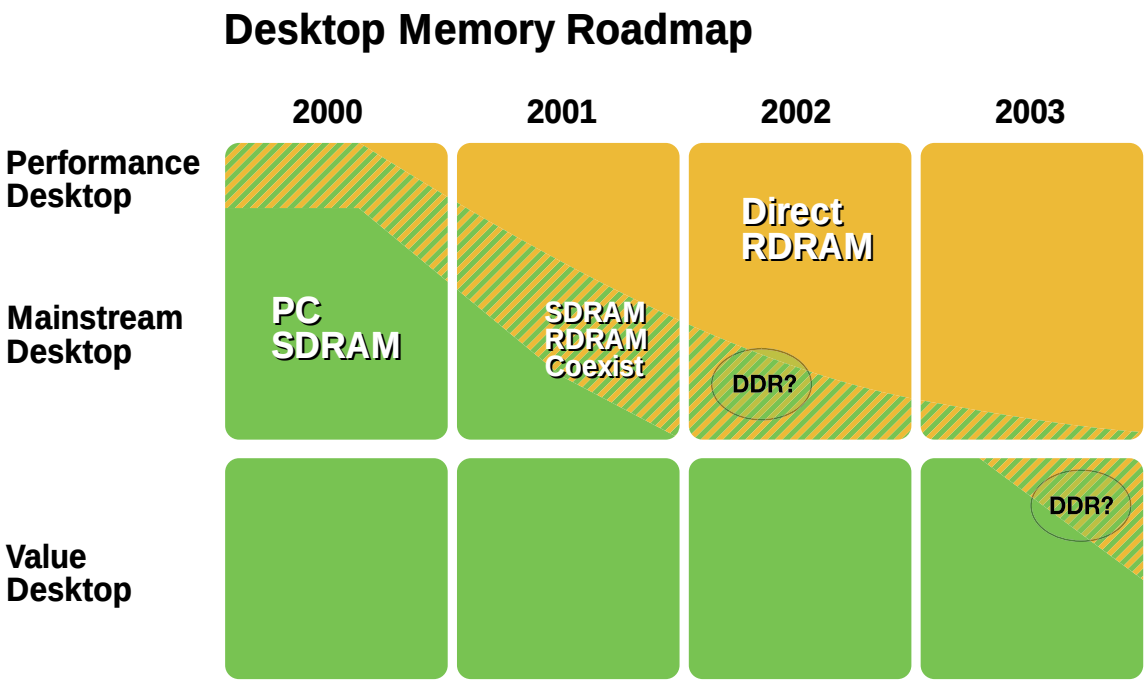

All dates are target dates and subject to change

- **SDRAM to RDRAM transition driven by Pentium® 4 Processor ramp**
- **SDRAM will continue to be used in cost sensitive segments as long as it remains lowest price**
- **DDR Solution targeted for 1H/02. Most likely role would be as SDRAM replacement**
- **Competitive RDRAM pricing minimizes the role of DDR**

The graph reflects the growing acceptance of RDRAM as the best choice for high-performance Intel Pentium 4 systems. The ramp up of Pentium 4 systems and higher volume production of RDRAM memory go hand in hand. As memory prices fall with higher volume production, system prices also drop. Pentium 4 processor-based systems are already beginning to penetrate the mainstream segment of the desktop marketplace. This year Intel will enable even lower price points for Pentium 4 systems by offering a new chipset that supports PC133 SDRAM memory. This will build on the extensive SDRAM support already offered by Pentium® III systems with the Intel® 815 chipset and the Intel® 810 chipset.

In 2002, Intel will introduce DDR memory support for the Pentium 4 processor through a new chipset. This new option will enable developers to bring Pentium 4 performance to an even wider range of price points. This will help developers in both the performance and mainstream market segments better respond to changes in memory price and availability.

RDRAM—The Choice for High Performance

RDRAM memory performs especially well with processors and systems that take full advantage of its exceptional memory bandwidth, such as the Intel Pentium 4 processor and its Netburst™ microarchitecture. Systems built on the Pentium 4 processor, the Intel 850 chipset, and RDRAM memory are providing users with the best available performance on the most demanding applications. The increasing acceptance of the Pentium 4 processor as the performance leader is driving higher volume production of RDRAM modules and significantly lowering RDRAM costs.

RDRAM technology has the advantage of being stable—RDRAM modules have been shipping for more than a year. Memory producers are completing the aggressive learning curve to achieve high-volume production. With yields increasing and prices dropping, Pentium 4 systems using RDRAM solutions are reaching lower price points sooner than expected.

Because of RDRAM's high speed, technological stability, and robust performance it has already become Intel's primary memory technology for high-performance market segments. Intel will continue to drive the development of platforms that take advantage of its higher bandwidth and growing availability.

DDR—Developing Technology

Intel is fully engaged with the industry to develop DDR memory technology. The Intel® DDR 200 spec addendum tightens and clarifies the Joint Electron Device Engineering Council (JEDEC) definitions to ensure that DDR components and dual in-line memory modules (DIMMs) will be as robust as possible. Intel is also focusing on motherboard designs to ensure that they are robust, cost-effective, and manufacturable.

DDR memory is already a suitable choice for certain applications. Accordingly, Intel is developing a DDR-based server platform using approved DIMMs for release later this year. As the volume production of DDR memory becomes more stable, Intel will expand DDR options for performance and mainstream desktop systems using more cost-effective unbuffered DIMMs. A new chipset that supports DDR memory will be available for Pentium 4 systems in early part of 2002.

PC133 SDRAM—The Value Solution

Currently SDRAM is the lowest risk, lowest cost, and most available memory technology. It allows developers to achieve reasonable performance while keeping system costs down. Beginning this fall, a new chipset will offer full support of PC133 SDRAM memory for Pentium 4 processor-based systems. This builds on the Pentium III processor's support of PC133 and PC100 SDRAM memory. Developers in the mainstream and value markets will have more options that allow them to more easily adapt to changing market conditions.

Summary

Intel is fully committed to ensuring that developers have a full range of platform memory solutions. The best choice for giving end users the finest computing experience will continue to be high-bandwidth RDRAM memory. Higher volume production and a stable technology will continue to lower unit costs and allow RDRAM-based systems to complete their penetration of the mainstream market segment. The development of DDR and SDRAM support for Pentium 4 systems will give developers additional options in addressing changing market conditions.

More Info

To find out more about Intel® memory initiatives, visit the Intel® Platform Memory page on the Intel Developer Web site.

To discover more about Pentium 4 processor performance, visit the Intel® Pentium® 4 Processor page.

For more information on the Intel 850 chipset, visit the Intel® Pentium® 4 Processor-based Platform with the Intel® 850 Chipset page.

Author Bio

Kuljit Bains is involved in Intel's strategic planning for memory technology. He has been with Intel for more than 10 years. He worked as platform architect for workstation and desktop systems and holds 10 patents in memory technology and related areas. Kuljit earned a B.S.E.E. from the University of Allahabad in India and a M.S.E.E. from California State University, Sacramento.

Serial ATA—the Long-term Solution for Storage Connectivity

Jeffrey Ravencraft
Strategic Marketing Manager
Intel Architecture Lab
Intel Corporation

Overview

Industry leaders recently announced that draft specification 1.0 defining Serial ATA (SATA), the high-performance interface for storage devices, is now available. This interface is used to connect storage devices such as hard disks, DVD, and CD-RW drives to the PC motherboard. Serial ATA delivers a scalable interface solution supporting several speed doublings to address the needs of several generations of future storage devices.

Serial ATA technology will replace today's parallel ATA physical storage interface. It enables the industry to move to the lower voltages and lower pin count required for efficient integration in future chipsets. What's more, it considerably improves the cable and connector plant, thereby improving ease of use and manufacturability. Serial ATA also provides opportunities for new capabilities not supported by parallel ATA, and it allows software to be streamlined by eliminating the master-slave interaction between devices.

Intel formed an industry working group for Serial ATA in February 2000. Promoter members of the group include APT Technologies, Dell Computers, IBM, Maxtor Corporation, Quantum Corporation, and Seagate Technologies. The specification reached draft 1.0 status by December 2000, with more than 70 contributing companies covering the disciplines required to enable this technology.

Significant Improvements over Parallel ATA

Today's parallel ATA interface has fundamental shortcomings that make it unsuitable as a long-term solution for future generations of computer technology. The Serial ATA specification addresses these shortcomings and provides a scalable platform to support several generations of future storage devices.

But transitions to new technology do not take place overnight—the process starts with discrete components and evolves over time. In order to meet the time-to-market opportunity for device deployment and subsequent high-volume ramp, manufacturers should prepare for the transition to the Serial ATA interface as the primary internal storage interface in early 2002. They should also ensure that they accommodate Serial ATA in their storage interface roadmaps and align accordingly.

The Serial ATA interface addresses several key problem areas:

Lower Voltage. Parallel ATA is based on TTL signaling, which requires integrated circuits to tolerate input signals as high as 5 volts. In the near future, integrated circuits manufactured on the leading manufacturing processes will not be able to efficiently support 5 volt signaling voltages. With the upcoming fine lithographies, it will not be feasible to continue supporting 5 volt signaling tolerance. Serial ATA addresses these integration issues by reducing the signaling voltages to approximately 250 millivolts (¼ volt).

Pin Efficiency. Currently, the parallel ATA interface has 26 signal pins going into the interface chip. Serial ATA uses only 4 signal pins, improving the pin efficiency and accommodating a highly integrated chip implementation.

Improved Cable and Connector Plant. The current parallel ATA cable and connector plant (Figure 1) is a bulky cable nest made up of unwieldy 80 conductor ribbon cables and 40 pin header connectors.

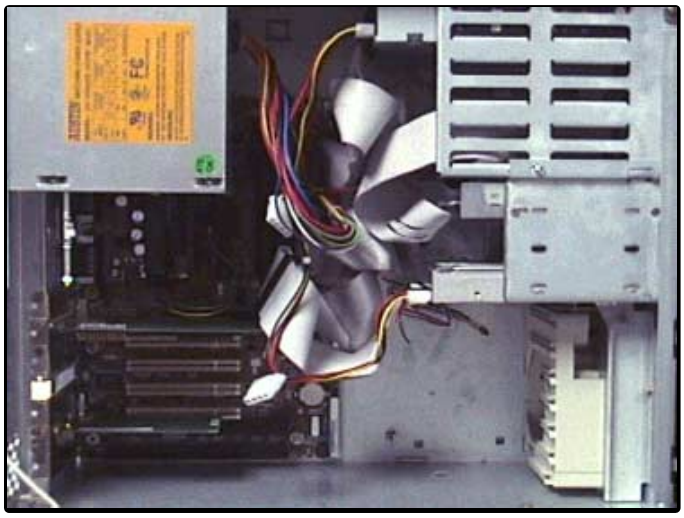

Figure 1. Parallel ATA Cable and Connector Plant

Instead of 80 conductor ribbon cables, Serial ATA (Figure 2) uses a much smaller serial cable similar in appearance to modern telephone cables. Eliminating the cable nest improves the system's airflow and cooling, and offers greater freedom in chassis design. It also improves manufacturability and ease of use by allowing better connector and cable design.

Figure 2. Serial ATA Cable and Connector Plant

Master-Slave Interaction. With today's parallel ATA implementation, pairs of devices share a common cable in a master-slave relationship. This interaction between devices results in the available bandwidth being shared between the devices. Additionally, since the devices on the cable interact, they must be jointly qualified, resulting in the substantial expansion of the system integrators' qualification matrix in order to comprehend all possible combinations of devices.

By contrast, Serial ATA is a point-to-point interface where each device is directly connected to the host via a dedicated link. Each device, therefore, has the entire interface bandwidth dedicated to it, and there is no interaction between devices. This means software can be streamlined, eliminating the overhead associated with coordinating accesses between the master and slave device sharing the same cable.

Hot-Plug Opportunity. An additional benefit of Serial ATA is that it provides the opportunity for devices to be hot-plugged and inserted directly into receptacles, an approach that is not directly supported by parallel ATA. Serial ATA includes all the mechanical and electrical features necessary to allow devices to be directly inserted into receptacles while the system is powered (commonly referred to as "hot-plugged"), and the protocol ensures that device discovery and initialization are handled.

100 Percent Software Compatible

Serial ATA is compatible at the register level with parallel ATA. This means Serial ATA requires no changes to existing software and operating systems in order to function, and it provides backward compatibility with existing operating environments. It has already been demonstrated to run with Microsoft Windows* 98, Microsoft Windows 2000, and Linux* operating systems without the need for any special software drivers. It is completely backward-compatible from a software standpoint.

Performance Headroom for the Next 10 Years

Parallel ATA does not have scalability to support several more speed doublings, and it is nearing its performance capacity. By contrast, Serial ATA defines a roadmap starting at 1.5 gigabits per second (equivalent to a data rate of 150 MB/s) and migrating to 3.0 gigabits per second (300 MB/s), then to 6.0 gigabits per second (600 MB/s). This roadmap supports up to 10 years of storage evolution, based on historical trends.

Summary

The Serial ATA specification now provides a stable platform for device development and deployment. The working group believes Serial ATA will deliver the best long-term storage interface solution, addressing the shortcomings of parallel ATA while delivering scalable performance to support an interface roadmap spanning at least 10 years. It enables the industry to move to the lower signaling voltages and reduced pin counts required for efficient integration in future chip design. At the same time, it considerably improves the cable and connector plant. Furthermore, Serial ATA should improve ease of use and manufacturability, and streamline qualification of software by eliminating master-slave interactions.

More Info

For the Serial ATA specification, press releases, and additional presentation collateral, visit the Serial ATA Working Group site.

Author Bio

Jeffrey Ravencraft joined Intel in 1988. Before working in Intel Architecture Lab, he was a purchasing manager for the Personal Computer Enhancement Operation and business development manager of Intel's networking business with Compaq Computers. He won two Intel Achievement Awards in 1989.

InfiniBand* Architecture-Enabling Technology

Jim Pappas
Director of Initiative Marketing
Fabric Component Division
Intel Corporation

Overview

InfiniBand* architecture is an industry initiative for high-speed, serial I/O, switched "fabric" architecture. It resolves many of the traditional barriers seen with shared-bus architectures, and has become the I/O interconnect for next-generation data-center communications. With broad industry support, the InfiniBand architecture is one of the key elements in simplifying and building reliable, low-cost, high-performance server systems.

After the Intel Developer Forum (IDF) Conference, Fall 2000, the InfiniBand Trade Association finalized the InfiniBand Specification 1.0. More than 45 companies have now announced product intentions for channel-based InfiniBand products. Many vendors, including Intel, have announced sampling of InfiniBand components for volume commercial delivery later this year.

At previous IDF conferences, Intel has shown working concept and prototype InfiniBand products. At the recent IDF Conference, Spring 2001, Intel demonstrated the first working InfiniBand fabric, offered InfiniBand enabling software and product development kits, and announced Intel's InfiniBand interoperability lab.

First Demo of Fabric

The demonstration of the first working InfiniBand fabric was based on Intel's recently announced InifiniBand silicon component. In the future, Intel's InfiniBand products will integrate into Intel® Architecture 32-bit (IA-32) and upcoming Intel® Itanium™ processor-based server platforms. InfiniBand lets developers take advantage of the performance, scalability, and reliability of the IA-based servers that are expected to be commercially released later this year.

The InfiniBand fabric demonstration showcased three protocols transferred through a single fabric and on the same wire (dual simplex), running at 2.5 Gbits/s. It comprised eight vendors' InfiniBand components and demonstrated three protocols (storage, networking, and clustering). The vendors that participated in the InfiniBand fabric demo were Adaptec, Agilent, Compaq Computer, Computer Associates, Crossroads Systems, Intel, LSI Logic, and QLogic.

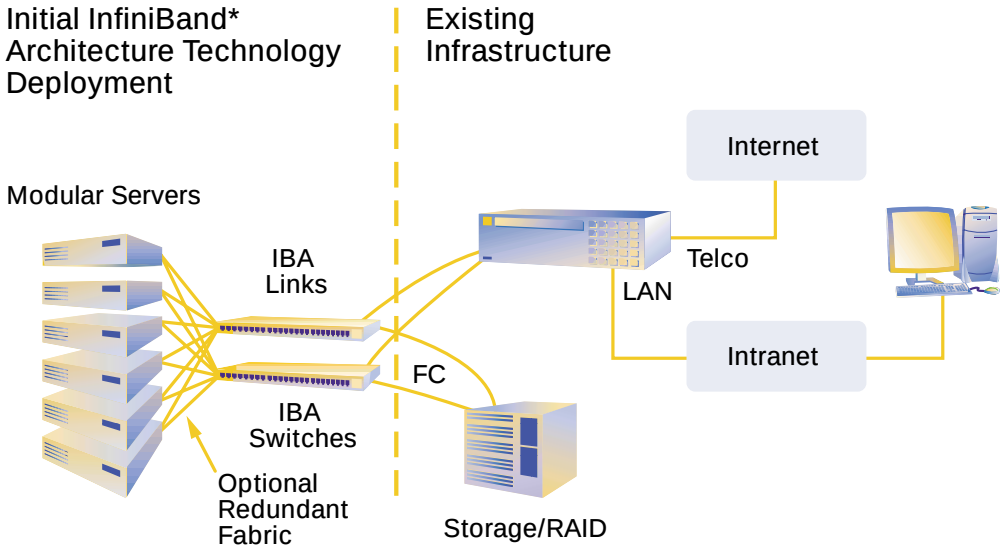

Figure 1. Initial InifiniBand* Architecture Technology Deployment

Enabling Programs

Parallel development helps third-party vendors ensure that their products are market-ready by the technology launch

date. To help vendors with parallel product development for InfiniBand servers, switches, peripherals, and other parts of the environment, Intel has created several enabling programs:

- Product development kits (PDKs), comprising specification-compatible components and software
- Port logic, which creates a common interface for InfiniBand components to help speed interoperability
- Target transport services (TTS), a suite of software components that help deliver interoperable platforms for discrete InfiniBand components

PDKs

Along with the requisite software and cables, the InfiniBand PDK has three critical components:

- *A host-channel adapter (HCA)*, which includes the transport services library, a connection manager, a channel-access layer, and drivers
- *A target channel adapter (TCA)*, which includes various drivers and which may also include an HCA stack
- *A switch*, which connects servers to remote storage and networking devices

Currently, PDKs support Microsoft Windows* 2000, Linux*, and VxWorks* operating systems.

Port Logic

Port logic is the fundamental software that lays the foundation for an InfiniBand software stack. Port logic drives the software stack at the link layer and helps ensure interoperability between InfiniBand products. Intel is making InfiniBand port logic widely available through licensing to multiple vendors. Contact your IT house or foundry for more information about obtaining Intel's InfiniBand port logic.

Target Transport Services

Target Transport Services (TTS) are a suite of software elements that help developers accelerate development of target adapters. For example, TTS helps make sure the software stack of the target channel adapter (TCA) is compatible with the software stack of the host channel adapter (HCA). TTS can also help ensure that the software stack is robust, and it can help verify interoperability with InfiniBand hardware.

Interoperability Lab

Since the finalization of the InfiniBand Specification 1.0, Intel has created the first InfiniBand interoperability lab for developers and vendors. This is a physical lab located in Oregon and available for the use of the InfiniBand industry. Developers can sign up for access to the lab to test the interoperability of InfiniBand products with Intel’s InfiniBand products. For more information about the permanent lab, contact your Intel representative.

Summary

InfiniBand architecture-enabling technology is available now, including product development kits, port logic, and target transport services. These enabling technologies give developers specification-compatible components and software, a common interface for InfiniBand components, and software suites for interoperable InfiniBand platforms. Intel has also created the first InfiniBand interoperability lab, which is available to developers across the InfiniBand industry.

With the finalization of the InfiniBand Specification 1.0 and the product intention announcements of over 40 InfiniBand vendors, developers should start taking advantage of the newly available enabling technologies.

More Info

For more information about Intel’s enabling programs, visit the Intel Developer Web site.

For information about obtaining Intel's InfiniBand port logic, contact your IT house or foundry.

For more information about testing the interoperability of products in Intel’s InfiniBand interoperability lab, contact your Intel representative.

Author Bio

Jim Pappas works with the industry on the development of InfiniBand products and promotes the standard through the InfiniBand Trade Association. He has served on the boards of numerous technology initiatives, most notably as a founding member of the PCI Special Interest Group and as founding chairperson of the USB Implementers Forum. Previously, Jim was director of Technology Initiatives in Intel's Desktop Products Group. During his 18 years in the computer industry, he has successfully advanced desktop technologies including AGP, DVD, IEEE 1394, Instantly Available PC, and USB. Jim holds eight U.S. patents in the areas of computer graphics and microprocessor technologies. His B.S.E.E. is from the University of Massachusetts in Amherst.

Intel Enhances the Intel® IXA Processor Family

Alex Shoykhet
Product Marketing Engineer
Edge Processing Operation
Intel Corporation

Overview

Intel recently announced two major enhancements to the Intel® Internet Exchange Architecture (IXA) family of network silicon and software building blocks. These enhancements are an evolution of the Intel® IXP1200 network chips and tools that offer three new pieces of technology: a new, faster network processor, an enhanced software development environment, and the new, high-level language Intel® Microengine C Compiler.

Both the Intel® IXP1200 Network Processor and the Intel® IXA Software Development Kit (SDK) version 1.2 for the IXP1200 are available now. The Microengine C Compiler will be available in spring 2001.

Microengine C

Previously, developers have used structured assembly language (IXP Symbolic Language) to program the IXP1200. However, with the announced enhancements, developers can now also program the IXP1200 using the high-level language Microengine C.

Microengine C represents a new enhancement to the programming options for the IXP1200 network processor. Microengine C includes special language additions called intrinsics. These additions, which are extensions to the standard ANSI C language, are optimized high-level abstractions for specific IXP1200 hardware functions. For example, these intrinsics handle CSR (control store register) unit programming, specific memory transfers, synchronization, and other tasks unique to the IXP1200. Microengine C intrinsics let developers concentrate on developing features for product differentiation rather than focus on nuts and bolts assembly-language functions.

Microengine C Compiler

The Microengine C Compiler is the first compiler for the IXP1200 Network Processor. The compiler is a strong addition to the IXP1200 Network Processor tool set. The Intel Microengine C Compiler supports all ANSI C statements and a subset of the ANSI C library and data types.

With the addition of this Microengine C compiler, developers can now choose to write prototype and production code in a high-level language, an assembly-level language, or a combination of both. For example, the compiler supports the in-line assembly feature, so developers can optimize specific areas of their Microengine C programs or reuse previously written IXP Symbolic Language code.

Because the Microengine C Compiler allows developers to program in a high-level language, it not only makes it easier to code today's network processors, it also provides a portability path for the next generation of processors.

Intel® IXA SDK Version 1.2

One of the IXP1200 software enhancements showcased at the Intel Developer Forum (IDF) Conference, Spring 2001 is Intel® IXA SDK version 1.2. This version of the Intel IXA SDK supports the latest version of the IXP1200 silicon and includes a library of more than 100 macros optimized for the IXP1200 network processor. These machine routines include macros for operating system emulation, instruction simplification, utilities, networking functions, and I/O operations. For example, instead of writing new code for common tasks, such as receiving packets from the IXP1200 I/O bus, developers can use a prewritten and optimized macro function. In combination with the Microengine C Compiler, the Intel IXA SDK v1.2 macro library helps speed up the development and optimization of value-add applications.

IXP1200 Network Processor

The Intel IXP1200 Network Processor includes an embedded Intel® StrongArm* control processor core and six microengines. The microengines operate on data-path packets to perform header parsing and look-ups, packet transformations, and other functions. Each of the six microengines has its own instruction control store and can execute four threads, with context swapping on memory references.

Now, Intel has announced the newest version of the IXP1200 network processor. New features include significantly increased core and bus speeds:

- Core clock speed has been increased from 200 MHz to 232 MHz.
- External memory bus speed has been increased from 100 MHz to 116 MHz.
- I/O bus speed has been increased from 85 MHz to 104 MHz.

The processor instruction store (the on-chip control store) has also been doubled, from 1K to 2K instructions per microengine. An industrial temperature-compliant version of the IXP1200 will be available in spring 2001.

The combined benefit of the IXP1200 enhancements is that developers can now execute individual instructions faster, and execute more instructions per packet in the data path. This means support of faster networking interface speeds and higher throughput, as well as greater instruction headroom for value-added network applications.

Summary

Intel announced a set of enhancements to the IXP1200 network processor programming toolset at the Intel Developer Forum Conference, Spring 2001. The Microengine C Compiler and Intel IXA SDK version 1.2 are major enhancements that will make it easier to quickly develop new networking applications and reduce product time-to-market.

More Info

To order the entire Intel IXA SDK version 1.2 developer's environment, visit the Intel Developer Web site.

Author Bio

Alex Shoykhet is currently with Intel as a senior product manager in the Edge Processing Operation group. He has managed the product development and integration of the Intel IXA SDK version 1.2 for the IXP1200 network processor, as well as the Microengine C compiler. Previously, he worked as a staff engineer at Netboost (acquired by Intel in 1999) and at Altera. Alex earned his B.S. and M.S.C.S. from Stanford University.

Innovative PCs Deliver Performance, Simplicity, Personality

Gabriel Achanzar
Innovative PC Program Manager
Consumer Desktop Marketing Group
Intel Corporation

Overview

At the recently concluded Intel Developer Forum (IDF), Spring 2001, seven PCs were recognized for their innovative approach to the longstanding industry challenge of delivering functionality, usability, and style in a single package. Showcasing the imagination and hard work of dozens of developers, the Intel® Innovative PC Recognition Program awards went to the Compaq* iPAQ in the business desktop category; the Fujitsu-Siemens* Jetson, Sony* Digital Studio, and Sotec* AFiNA AV in the consumer desktop category; and the Compaq Presario 800, Dell Latitude C800, and Samsung* Sens NV5000(SENS 760) in the mobile category. To view the winning designs, click on the hyperlinked winner's individual names.

Compaq* iPaq

Fujitsu-Siemens* Jetson

Sony* Digital Studio

Sotec* AFiNA AV

Compaq* Presario 800

Dell* Latitude C800

Samsung* Sens QS760

Of the entries from developers around the globe, these seven best represented the Innovative PC Recognition Program's theme of Performance, Personality, and Simplicity, and they did it by delivering a set of features that improve the user experience.

PC and Home Entertainment Center Rolled Into One?

In the area of performance and personality, for example, many of these machines are helping to further blur the line between what users traditionally think of as PCs and devices for the production and delivery of entertainment. Consider video and audio capabilities for starters:

- The Compaq Presario 800, Dell Latitude C800, and Sony Digital Studio, Sotec AFiNA AV feature an IEEE 1394 port enabling users to capture signals from a digital video camera for editing and storage on the PC.
- The Sony Digital Studio literally packs an audio, video, and photo editing studio into a powerful PC. Equipped with a powerful Pentium® 4 1.5-GHz processor, the Digital Studio has the headroom for these demanding tasks.
- The Dell Latitude C800 and Sotec AFiNA AV provide an S-Video output so that a DVD playing on the PC can have its display signal sent to a large-screen television.
- The Sotec AFiNA AV includes an RCA port as well as a remote control for both its internal DVD and for a television tuner, for the ultimate flexibility in receiving and displaying video input. The Sotec AFiNA AV also provides Optical Audio, which enables users to send an audio signal over a high-speed, high-quality fiber-optic connection to traditional audio device such as a stereo.

Easier From the Start

As for simplicity, all seven PCs feature an impressive selection of capabilities designed to significantly ease the work of setup, operation, and expansion:

- All the models feature USB, the serial connection technology that supports the addition of peripherals without requiring any hardware installation. The Compaq iPAQ, Fujitsu-Siemens Jetson, Sotec AFiNA AV, and Sony Digital Studio feature the USB ports on the front panel so even just plugging in a peripheral becomes easier. Moreover, the Compaq iPAQ, free of older ports such as the serial, parallel, and game ports, came with an innovative adapter for connecting traditional peripherals.
- The Compaq Presario 800, Dell Latitude C800, Samsung NV5000, and Sotec AFiNA AV include a PCMCIA slot for easy addition of such enhancements as Bluetooth* wireless capabilities or additional memory.
- Unusual among mobile computers, the Dell Latitude C800 features both DVD and CD-RW, and the DVD is removable so its slot can be used for a spare battery or floppy drive. It also features an internal antenna making ready for wireless connectivity. Likewise, the SOTEC AFiNA AV comes with a DVD/CDRW occupying a single bay.

- The built-in IEEE 1394 capabilities of the Compaq Presario 800, Dell Latitude C800, and Sony Digital Studio are not only a performance advantage but also an ease-of-use advantage. That's because they enable users to input video without having to install additional software or implement complicated synchronization techniques between the input device and the PC. They also eliminate concerns about compression because with IEEE 1394 the hardware automatically uses a standard high-quality compression rate.
- For a versatile approach to portability, the Samsung NV5000 mobile PC comes with an equally mobile docking station weighing in at just two pounds. Both the DVD player and two of the USB ports are on the docking station while one more USB port is on the machine itself.
- At just an inch wide and less than three-and-a-half pounds in weight, the Compaq Presario 800, too, provides excellent portability. For an easier way to listen to audio, this mobile PC also features a front-panel volume adjuster, just like a radio or CD player, in contrast to the traditional mobile-PC approach of requiring users to go the keyboard.
- For making two of the most common PC tasks a lot faster and easier for users, the Compaq Presario 800 and Samsung SENS NV5000 feature one-touch Internet and e-mail access.
- Often thought of as an energy-saving feature but one that also makes life easier for PC users is IAPC (Instantly Available PC), a technology that implements a very low-wattage standby mode and makes turning a PC on as simple—and fast—as turning on a radio. The Compaq iPAQ feature IAPC.

And Looks Good, Too

Personality, the third criterion in the selection of this year's IDF winners, was represented in both design features and capabilities:

- The Samsung SENS NV5000 comes with an optional miniature video camera and an MP3 player, both with an exterior design matching that of the PC which comes in a sleek metallic gold case.
- The Fujitsu-Siemens Jetson, in a compact silver case that looks like something James Bond might carry, is attractive on the desktop in either a horizontal or vertical configuration.
- The Sony Digital Studio also features a smaller than usual footprint thanks to its use of the Intel® microATX motherboard, which manages to comfortably house a Pentium 4 processor.
- The Sotec AFiNA AV, in a sleek design featuring an aluminum finish with metallic blue trim, looks at home sharing a shelf with the most sophisticated of stereo equipment.

Summary

The features and capabilities represented by this year's Innovative PC Recognition Program's winners are described here according to the forum's theme of Performance, Personality, and Simplicity. But they're best thought of as an integrated and deliberate effort on the part of dozens of designers and developers to give users a more productive and rewarding PC experience.

With the help of enabling technologies from Intel such as IAPC technology, and the FlexATX and microATX motherboards, and with technology initiatives such as USB, and DVI, these innovative PCs are giving users performance capabilities until now available only through production studios equipped with high-end equipment for producing and enjoying video and audio. They're making it far easier than ever for users to handle setup, daily operation, and enhancements. They're attractive to look at and they fit easily in even the most crowded of home or office surroundings.

More Info

To maintain the momentum of PC development efforts like these, Intel works closely with its associates in the Ease of Use Roundtable and through endeavors such as the Ease of Use Initiative and the Innovative PC Recognition Program.

For more information, visit the Ease of Use Roundtable Web site or the Intel Developer Web site. To view the Innovative PCs, please visit the Intel Developer Web site.

Additional information is available through the author of this article, Gabriel.r.Achanzar@intel.com.

Author Bio

Before joining Intel in 2000, Gabriel Achanzar worked in the semiconductor equipment industry. Prior to his role as Innovative PC Program Manager, Gabriel provided marketing strategy and support for the interactive technologies for the Dave Matthews Band Nationwide tour, the Windows Me Launch, and Ease of Use exhibit at the Intel Museum. He holds an M.B.A. in technology management from the University of Phoenix.

Servers

Cross-Platform Remote Server Management Capabilities

Tom Slaight
Server Management Architect
Enterprise Platforms & Services Division
Intel Corporation

Overview

Version 1.5 of the intelligent platform management interface (IPMI) specification takes a crucial step in resolving one of the major challenges of server management. It introduces the common interfaces needed to develop cross-platform remote platform management software.

By incorporating LAN and serial/modem access, monitoring, alerts, and automated emergency responses, this new version of the IPMI builds on the proven "always available" management technology first introduced by IPMI v1.0 in September 1998. Now software developers can build in remote management features that operate efficiently across system types while gaining increased speed, efficiency, and control over the process of development (see Figure 1).

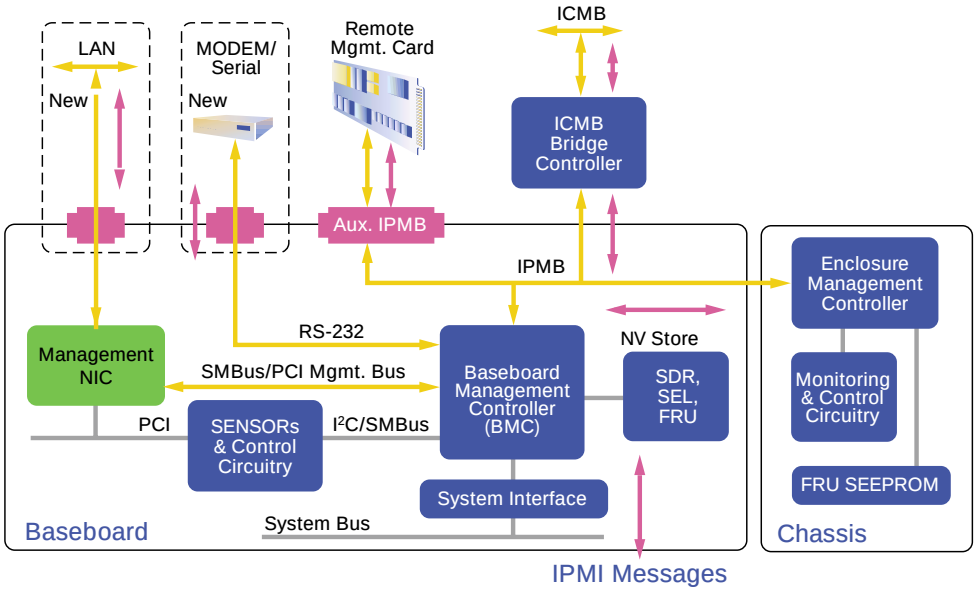

Figure 1. IPMI v1.5 adds new LAN and serial/modem interfaces to support remote management capabilities.

New Remote Access Options Meet User Needs

IPMI v1.0 offered just two mechanisms for remote access: via a remote management add-in card connected to an auxiliary intelligent platform management bus (IPMB) connector, and from other systems via an intelligent chassis management bus (ICMB) connection. With version 1.5 LAN and serial remote, capabilities can be built directly into the management system, giving users a more economical way to check on the status of their system or handle routine maintenance and inventory operations.

Specific capabilities include: monitoring system health information, power on and off, and recovery operations, such as hard reset or diagnostic interrupt. Remote administrators can also set "boot options" that direct the way the system boots after a reset or power up. For example, a system could be directed to boot to a diagnostic "service" partition instead of the main OS partition.

IPMI provides “symmetric manageability.” A common set of management capabilities is supported over all of the specified interfaces, so users can choose the access method that best meets their operating needs. This also saves the user and the software developer from the confusion and frustration that can result when capabilities are available through one interface but not another.

Users, Privileges, and Authentication

It is always important to provide authenticated access to systems, particularly on interfaces that allow a remote administrator to perform operations such as system resets. The LAN and serial/modem interfaces include user authentication mechanisms that support multiple users and privilege levels, and users can be separately enabled with different privileges on each interface. This capability can be used to set up one set of users for modem access and a different set for LAN access.

Links with Other Standards and Initiatives

Version 1.5 takes advantage of and supports several related existing and emerging platform management standards and initiatives, including:

- PPP for serial communications (including support for standard PPP authentication protocols such as CHAP, PAP, and MS-CHAP)
- MD2 and MD5-based authentication for serial and LAN interfaces
- DMTF (pre-OS working group) alert standard forum specification
- SMBus 2.0
- Compact PCI
- The proposed PCI management bus

By supporting both widely used protocols and emerging standards, IPMI v1.5 enables developers to take advantage of existing knowledge and tools while advancing to support new management technologies.

Serial Port Sharing

Serial/modem access is further enhanced by serial port sharing. This specification defines interfaces that enable the motherboard and management micro-controller to share access to a single serial connector on the back of the box. This approach enables a remote application to alternately communicate with the management controller and system or BIOS software without requiring multiple connections or modems.

Serial port sharing can be used in conjunction with the serial console redirection capability that many systems offer, allowing operations such as remote BIOS and OS configuration, while retaining the ability to access the management controller for system health status and to initiate remote system power on/off and reset actions. The definition of serial port sharing has been designed to be compatible with the Microsoft “Whistler Headless” specification for service controller access.

Intelligent Automated Recovery Features

IPMI v1.5 also supports sophisticated recovery features, using a mechanism called platform event filtering (PEF). Whenever a system event occurs, it is compared to configurable set of event filters. If there’s a match, an automated response is triggered. Specific actions include power off, power cycle, reset, diagnostic interrupt, and alert.

For greater control, these steps can be used in combination, for example to turn off the system when it overheats, and then send an alert to the system administrator. In this case, immediate action always takes priority over alerts. In addition, a “wild card” setting allows the filter to be configured very specifically—or more generically. For example, one filter might be set to match all “over-temperature” events, while another might be configured to match only “chassis intrusion on door 1” events.

Customized Alerts

With IPMI v1.5, system alerts can be provided over both LAN and serial/modem interfaces. Serial/modem alerting is supported as either numeric "dial" paging or using alphanumeric paging. LAN alerting is handled by generating SNMP traps following the platform event trap (PET) format specification, with optional support for confirming trap delivery. Serial/modem and LAN alerts are launched by PEF whenever an event matches a corresponding PEF entry that has "alert" as a selected action.

Both LAN and serial/modem alerts can be delivered to multiple destinations, where delivery can be customized with one or more "alert policies." For example, alerts can be sent to more than one LAN address, or more than one number can be dialed up.

In addition, destination types can be mixed and clustered in a configurable "call-down" policy, which might send an alert to a LAN, followed by a serial call, followed by another LAN destination. Success criteria can be associated with each step as a condition for the alert. For example, an alert can be set up to be initiated only if the previous call fails.

Multiple alerts policies can be configured. For example, a "high priority" alert might call a central service manager, while a "low priority" notifies the field service engineer. Each entry in PEF can select an alert policy, so the high-priority policy can then be associated with high-priority events, and the low-priority policy with low-priority events.

Advantages for the Development Process

IPMI v1.0 first enabled cross-platform management capabilities by establishing a set of specifications that define common, abstracted, message-based interfaces to intelligent platform management hardware. This approach allowed capabilities to be delivered across system types, while simultaneously enabling a faster, more efficient, more economical development process.

Specifically, the common interfaces of the IPMI reduce development efforts while still supporting the ability to differentiate with competitive new features. Abstraction provides further efficiency because it isolates software development from hardware development, allowing each to proceed independently without the delays and complications produced by synchronizing efforts. Building on this proven technology, Version 1.5 extends IPMI benefits to developers who want to create new remote management capabilities.

In addition, the latest version offers the valuable ability to efficiently reuse local software for remote management capabilities. By using the same abstractions and messages on the serial/LAN interfaces that were defined for the local interfaces, a "local" software stack could be converted to a "remote" stack just by changing the underlying communication driver used to access the system.

Benefits to Software Developers

Version 1.5 of the IPMI provides software developers these benefits:

- Remote management software can now operate across all platforms.
- Quality is improved. IPMI's common message-based interfaces work well with automated testing.
- Knowledge can be preserved and re-used.
- Common messaging across interfaces eliminates the need for experts in each interface; developers can easily move expertise across interfaces.
- Time-to-market is faster.
- Costs are reduced.
- Competitive new features are supported.

Benefits to Hardware Developers
Version 1.5 of the IPMI offers hardware developers these benefits:

- New features and innovations can be developed without worrying about synchronization with software.
- IPMI's interfaces allow a wide variety of components to be used for implementing the specification. Designers have freedom to implement new features and customize implementations to fit their system needs.
- Multiple component vendors offer management controllers, firmware, and development kits for creating IPMI-based systems.
- Time-to-market is faster.
- Costs are reduced.

Summary

IPMI v1.5 offers crucial advantages for developing cross-platform remote server management capabilities. Building on the strengths of IPMI v1.0, it extends the principle of "always available" manageability for users while providing developers with increased efficiency and faster time-to-market.

Although proprietary serial and LAN remote management capabilities have been available to users, they are typically more costly to implement and support across multiple platforms, and cannot readily take advantage of the new components and tools that come out when a common specification is available. IPMI v1.5 introduces the first common interfaces for remote server platform management, providing a more economical way to deliver monitoring, customizable alerts, and automated emergency responses.

These capabilities give users extensive control and a range of operating options to meet their needs. A choice of interfaces is supported, including standard LAN and serial/modem interfaces, as well as connections based on a growing range of emerging standards.

For developers, the latest version of IPMI expands the proven benefits of version 1.0 when it comes to increasing the speed, economy, and power of the development process. For this reason, Intel, HP, NEC, and Dell are all promoters of the IPMI specification, along with a growing group of over 70 IPMI Adopters representing a broad cross section of the server and server component industry. In addition, IPMI technology is called out in other industry specifications, such as Compact PCI and the DMTF alert standard forum specification. As a result, it's important for all developers to learn more about IPMI and how it can be applied in products.

More Info

Version 1.5 of the IPMI specification is available from the IPMI Web site. This site also contains additional IPMI information for developers, plus information on how companies can sign-up as an IPMI adopter.

You may also want to find out about third-party developers that make microcontrollers and firmware deploying IPMI, such as National Semiconductor, QLogic, Vitesse Semiconductor, and WinBond Corporation.

Author Bio

Tom Slaight has more than 15 years experience as a lead electronic design engineer, developing new and enhanced product architecture. He is an originator of the Intelligent Platform Management architecture and is a lead author and technical contributor to the IPMI specifications. Tom has participated in numerous other system management initiatives, receiving an Intel Achievement Award for his work on the InfiniBand* hardware management specification. Tom holds a B.S. degree in computer systems engineering from the University of Massachusetts, Amherst, and a M.S. degree in electronic engineering from Rutgers University.

Software

Using SSE and SSE2: Misconceptions and Reality

Alex Klimovitski
Software Application Engineer
Tools and Technologies, Europe
Intel GmbH

Overview

Single Instruction Multiple Data (SIMD) has industry-wide acceptance as a technology that increase performance on a wide variety of applications. Intel® Streaming SIMD Extensions (SSE) are sets of instructions that Intel introduced with the Intel® Pentium® III processor. This technology was the first major expansion of the Intel® instruction set since the introduction of Intel® MMX™ technology in 1995. With the launch of the Intel® Pentium® 4 platform, SSE was supplemented by SSE2. SSE/SSE2 can dramatically enhance software performance in a wide range of applications, from secure communication to speech recognition and synthesis, and from 3D visualization and video processing to realistic physics modeling.

Programming with SSE/SSE2 is quick, easy, and supported by mature tools. However, when trying to use SIMD in a real-life application, a developer is frequently confronted with a number of tough questions. How do I best exploit the parallelism promised by SIMD? How should I organize my data? And what if the optimal data structure isn't compatible with my algorithm? This article seeks to rectify some misconceptions about using SSE/SSE2 in the "real life" applications and to direct you to techniques and tricks that help achieve the best performance in your code.

SSE/SSE2 and Data Types

Misconception: Intel® SSE/SSE2 technology is applicable only to certain data types (just single precision floating-point, or just 16-bit integer…)
Reality: Intel SSE/SSE2 technology applies to elements of *all* standard data types, both floating-point (FP) and integer, that fit into 16-byte-wide SSE/SSE2 registers (Figure 1).

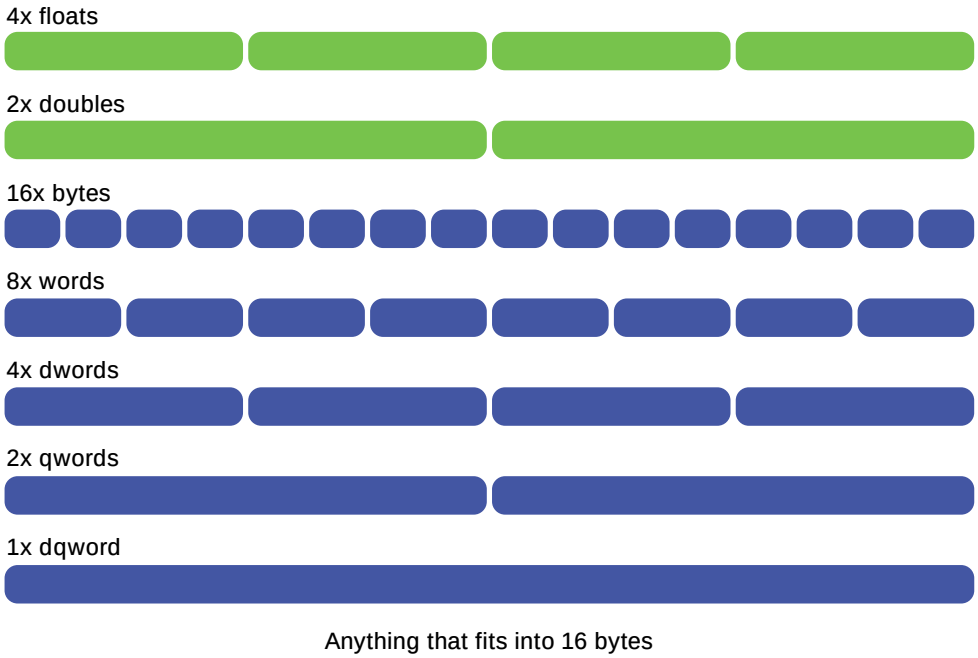

Figure 1. SSE/SSE2 data types

The supported data types include double- and single-precision FP data and all integer types with a length from 8 up to 128 bits. SSE/SSE2 offer a rich set of operations for each of these data types. This means, for example, that with just one instruction you can simultaneously add or multiply two operands of four floats, or of two doubles, or of two 64-bit integers, or of sixteen 8-bit integers, etc. This is how SSE/SSE2 dramatically increase the computational power available to your applications.

SSE/SSE2 and Coding Effort

Misconception: SSE/SSE2 require profound changes to the algorithm and significant additional coding effort.
Reality: Applying SSE/SSE2 lets you maintain the original algorithm. With SSE/SSE2, it can now process several independent data portions at once instead of one at a time. For example, for float data type, four data portions can be processed at once.

If you use C++ Vector Classes, you can also keep your original source code—just change the original scalar data type to the corresponding Vector Class that contains several elements of the type. For example, use F32vec4 Vector Class to keep and process four float elements in one operand.

The key is to keep only homogeneous elements in an SSE/SSE2 operand. For example, when dealing with an array of four-component (x, y, z, w) vectors, resist the temptation to stick all components of a vector into an SSE/SSE2 operand. Instead, work on the four vectors in parallel, and place their four x components in one SSE/SSE2 operand, the four y components into another, as well as four z and w components. Refer to Figure 2 to see how an array of four-component vectors is multiplied by a 4x4 matrix with perspective correction (x, y, and z components of the resulting vector is divided by w).

```
for (int i = 0; i < ARRAY_COUNT; i += 4) {
  F32vec4 x = (F32vec4&)xi[i];
  F32vec4 y = (F32vec4&)yi[i];
  F32vec4 z = (F32vec4&)zi[i];
  F32vec4 w = (F32vec4&)wi[i];

  F32vec4 wr = rcp_nr(x * q[3][0] + y * q[3][1] +
    z * q[3][2] + w * q[3][3]);

  (F32vec4&)xo[i] = wr * (x * q[0][0] + y * q[0][1]
    + z * q[0][2] + w * q[0][3]);
  (F32vec4&)yo[i] = wr * (x * q[1][0] + y * q[1][1]
    + z * q[1][2] + w * q[1][3]);
  (F32vec4&)zo[i] = wr * (x * q[2][0] + y * q[2][1]
    + z * q[2][2] + w * q[2][3]);
  (F32vec4&)wo[i] = wr;
}
```

Figure 2. Four-component vectors multiplied by 4x4 matrix using SSE/SSE2 with vector classes

This approach assumes that you are using SIMD-friendly Structure of Arrays (SoA) or “Hybrid” structures for your data, as shown in Figure 3. These structures allow quick and convenient loading or storing of data elements, e.g. four floats, to/from an SSE/SSE2 operand.

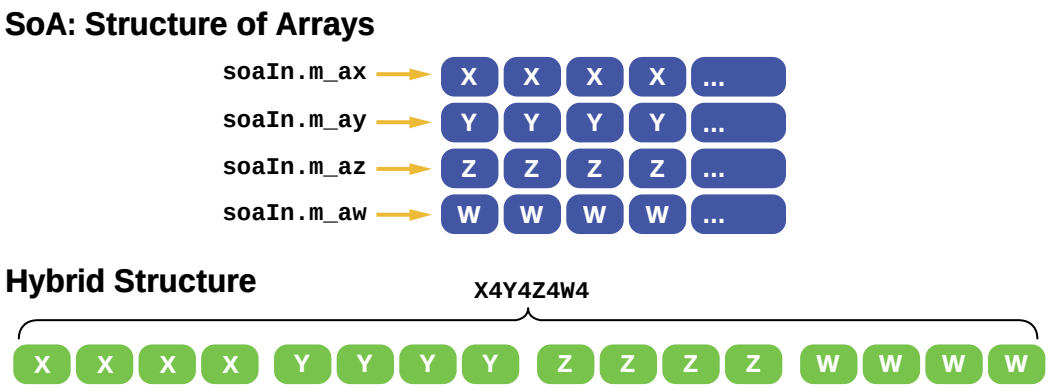

Figure 3. AoS (Array of Structures) hinders SSE/SSE2

SSE/SSE2 and Data Structures

Misconception: Because of interface/API constraints, algorithm logic, or legacy code, one still has to deal with SSE/SSE2-unfriendly data structures such as array of structures (AoS, Figure 1). These prevent any efficient application of SSE/SSE2.

Reality: It is most efficient to transform the data structure into SSE/SSE2-friendly form at design or load time. But if this is impossible due to interface/API constraints, algorithm logic, or legacy code, you can quickly modify the data structure on the fly by using SSE/SSE2. SSE/SSE2 provide a rich set of data transform operations such as half-loads, shuffles, packs, and unpacks. These help to load SSE/SSE2-unfriendly data and quickly "swizzle" it into a form suitable for fast processing with SSE/SSE2. If necessary, SSE/SSE2 also helps "unswizzle" the resulting data into the original form. As an example, Figure 4 shows an algorithm to transform AoS data into an SSE/SSE2-friendly form using SSE/SSE2 data restructuring operations.

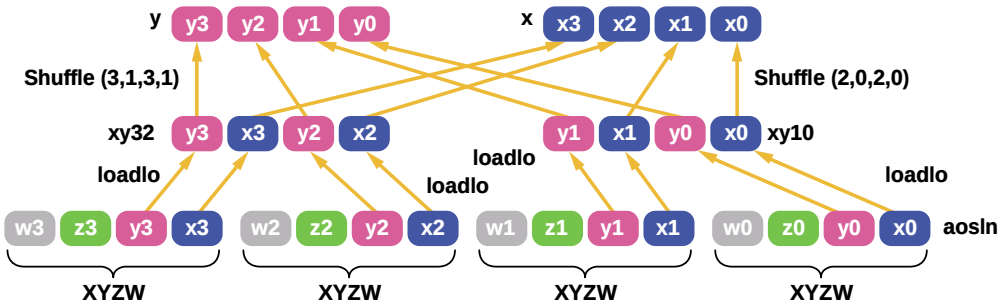


```
void XYZWtoF32vec4(F32vec4& x, y, z, w, XYZW* aosIn
{
F32vec4 xy10, xy32, zw10, zw32;
xy10 = zw10 = _mm_setzero_ps();
xy10 = _mm_loadl_pi(xy10, (__m64*)&(aosIn[0]).x);
zw10 = _mm_loadl_pi(zw10, (__m64*)&(aosIn[0]).z);
xy10 = _mm_loadh_pi(xy10, (__m64*)&(aosIn[1]).x);
zw10 = _mm_loadh_pi(zw10, (__m64*)&(aosIn[1]).z);
xy32 = zw32 = _mm_setzero_ps();
xy32 = _mm_loadl_pi(xy32, (__m64*)&(aosIn[2]).x);
zw32 = _mm_loadl_pi(zw32, (__m64*)&(aosIn[2]).z);
xy32 = _mm_loadh_pi(xy32, (__m64*)&(aosIn[3]).x);
zw32 = _mm_loadh_pi(zw32, (__m64*)&(aosIn[3]).z);
x = _mm_shuffle_ps(xy10, xy32, SHUFFLE(2,0,2,0));
y = _mm_shuffle_ps(xy10, xy32, SHUFFLE(3,1,3,1));
z = _mm_shuffle_ps(zw10, zw32, SHUFFLE(2,0,2,0));
w = _mm_shuffle_ps(zw10, zw32, SHUFFLE(3,1,3,1));
```

Figure 4. Transforming data from array of structures achieved with similar sequence of operations

SSE/SSE2 and Conditional Code

Misconception: Algorithms containing conditions are inherently unsuitable for SSE/SSE2. With SSE/SSE2 several data portions are processed, so many different combinations of conditions should be considered. This can be unpractical to implement and inefficient to execute.
Reality: You can successfully apply SSE/SSE2 even to conditional, branchy code and get a significant performance increase. The general approach is to replace the conditional branches with logical operations or computations.

Instead of setting the processor's condition flags as scalar comparisons do, SSE/SSE2 comparisons generate bit masks. All bits of the elements where the comparison yielded true are set to 1. All bits of the elements where the comparison produced false are cleared to 0. The mask can then be processed with SSE/SSE2 bitwise logical operations.

Figure 5 illustrates a general case of selecting the result depending on a comparison. Notice that several selections are performed in parallel—without a single branch. Since conditional branches often reduce code performance, the branchless SSE/SSE2 solution offers a significant performance benefit.

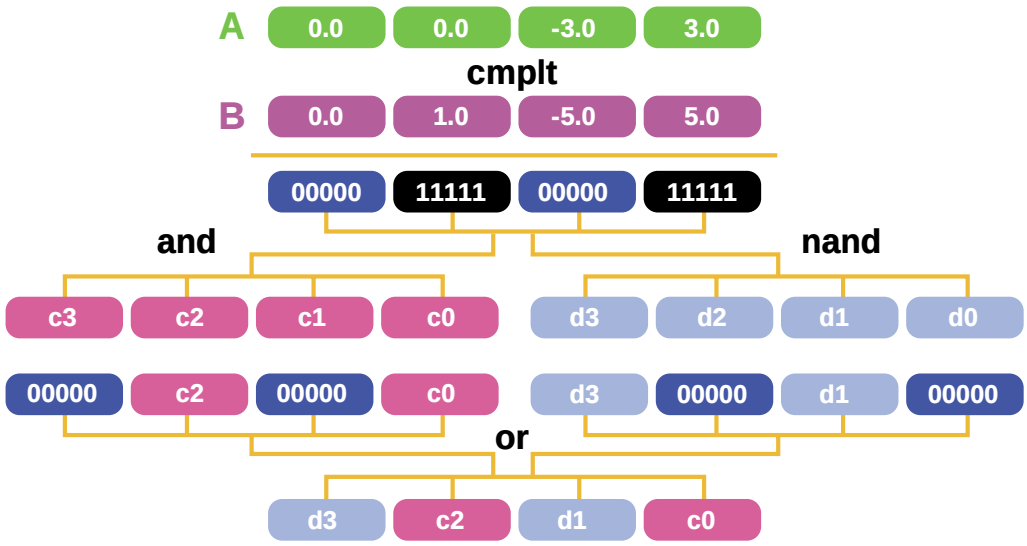


```
// R = (A < B)? C : D

F32vec4 mask = cmplt(a, b);
r = (mask & c) | _mm_nand_ps(mask, d);

// OR, using F32vec4 friend function:
r = select_lt(a, b, c, d);
```

Figure 5. Selecting the result depending on a particular condition

If a special reaction is required for a certain combination of conditions, you can apply a movemask operation to the bit mask produced by the comparison. Movemask generates the hash value of the mask by taking the most significant bit of every element. For a four-element comparison (e.g., for four floats of 32-bit ints), a value in the range 0000b…1111b will be produced. If, for example, a certain code block can be skipped because all element comparisons yielded false, then a single conditional branch does the job: the movemask value is compared against zero and, if equal, it invokes a branch around the block.

Summary
Programming with SSE/SSE2 is quick and easy. It uses mature tools and greatly improves performance of a wide variety of applications. Virtually all critical code can be accelerated with SSE/SSE2 processing. SSE/SSE2 can help transform SIMD-unfriendly data structures on the fly. And SSE/SSE2 comparisons and logic help eliminate branches in conditional code.

More Info
For more information on SIMD, SSE, and SSE2, visit the Intel Developer Web site.

Try out the Intel Compiler Interactive Tutorials in the software products area of the Intel Developer site.

For more information about developing with SSE/SSE2 visit the Intel Developer Services Web site.

Author Bio
Alex Klimovitski has been with Intel GmbH for five years. He develops tools and libraries for the latest Intel® Architecture processors, and assists leading software vendors in porting and enhancing their software for Intel® processors. He shares his experience in numerous presentations, application notes, and training courses. He holds an M.S. in computer science from Berlin Polytechnic.

Volume II of Intel® Developer Solutions Catalog Now Available

Will Swope
Vice President, Intel Architecture Group
and General Manager, Solutions Enabling Group
Intel Corporation

Overview

Volume II of the *Intel® Developer Solutions Catalog* was launched at the Intel® Developer Forum (IDF) Conference, Spring 2001. A compilation of third-party software tools and applications optimized or ported to Intel® Architecture, this catalog, available in printed, CD, and electronic formats, builds on the positive feedback generated by Volume I. Volume II incorporates many design improvements that make it easier to access and search listings.

The *Intel Developer Solutions Catalog* has also been considerably expanded; the new volume lists 50 percent more select independent software vendors (ISVs) than its predecessor. This one-stop solution reference is an expression of Intel's long-standing commitment to performance, manageability, flexibility, and openness in the developer community.

Benefits for Developers and Solutions Providers

The *Intel Developer Solutions Catalog* provides a number of key benefits:

- *Reduced time-to-market by streamlining code development.*
  Developers can reduce time and save money by building on tools and technologies that have already been tried and tested on Intel Architecture.
- *Deployment on multiple platforms and operating systems.*
  Developers can use the catalog to quickly locate solutions for a broad range of operating systems that run on Intel® processor-based platforms. This simplifies multi-platform development by minimizing the number of tools needed, and results in deployment on multiple platforms sooner.
- *Optimized performance and scalability.*
  Developers can create solutions that have been optimized and tuned for peak performance on Intel Architecture, achieving mission-critical dependability for high capacity demand and bandwidth.

New Design for Ease of Use

The catalog's new design offers several improvements:

- More specific categories help developers quickly find the types of software and solutions that interest them. A comprehensive functional index narrows searches. Each tool is listed on a separate page and arranged alphabetically so solution providers can be quickly located by name.
- Catalog concisely lists company, product overview, key features, and how the product showcases Intel Architecture.
- Comprehensive, CD-based format provides instant search results.
- Small, convenient size makes the printed catalog easy to carry and keep within reach.

More Resources for Developers

This catalog is only one of many resources available to software developers and solutions providers through Intel Developer Services. The IDS program provides a host of capabilities that enable developers to build solutions on Intel Architecture, including matchmaking and co-marketing opportunities, early access to the latest processor environments for application porting, training, software products and downloads, and a reference library. The IDS Web site includes an electronic version of the catalog, and includes links to the solutions listed, providing access to the most current data.

Summary
Volume II of the *Intel Developer Solutions Catalog* is now available. This catalog is a concise reference tool to help software developers quickly find applications optimized for Intel Architecture. It lists approximately 50 percent more ISVs than the previous volume, and has been redesigned for ease of use.

More Info
For information on how to order the *Intel Developer Solutions Catalog*, visit the Intel Developer Services Web site.

Author Bio
Will Swope joined Intel in 1979 as product marketing manager in the Development Systems division. He worked as director of strategic planning and systems operations, and marketing director for the Intel® Pentium® Pro and IntelDX2™ processors. In 1992, Will was promoted to co-general manager for the Pentium Pro processor division, and was named director for microprocessors in 1994. In 1996, he was promoted to vice president of Intel's Desktop Products Group, responsible for marketing Intel's complete product line of business clients.

Wireless

Building Blocks of Wireless Internet Clients

Vishwas Deshmane
Director, Group Communications and Marketing Programs
Wireless Communications and Computing Group
Intel Corporation

Overview

As today's mobile users demand more functionality and performance from wireless clients, developers are recognizing the tremendous opportunities that exist for hardware and software solutions that can support sophisticated Internet capabilities and meet user expectations. But they also face tremendous challenges. How can new products reach the market at the pace the public demands? How can mobile, battery-powered devices meet the standards established by users' desktop Internet experiences?

The Intel® Personal Internet Client Architecture (Intel® PCA) enables developers to meet the full range of market challenges. This blueprint for development includes three main building blocks that work together—but allow their development to occur in parallel—in order to enhance product performance, enrich functionality, and speed time-to-market.

The Intel PCA provides these advantages because, rather than merely extending the capabilities of old platform designs, it offers a new system architecture with a far more efficient development process. This groundbreaking approach to development significantly reduces programming efforts and compresses the time required to deploy software and hardware solutions.

An Approach that Accelerates Development

To keep pace with today's emerging wireless Internet market, Intel is addressing the development process in a whole new way. The Intel PCA separates the computing, communication functions, and memory, and links these three main blocks with a newly defined simple standard interface.

This approach allows hardware and applications development to proceed separately, in parallel, with applications written to a general-purpose applications processor as part of the computing block. When using a parallel development process, hardware and applications can evolve at their own speeds, but the overall time required to launch new products and services is greatly reduced.

This advance in efficiency takes a big step beyond today's slow, cumbersome serial development process. Figure 1 shows how the three main building blocks are linked by the interface.

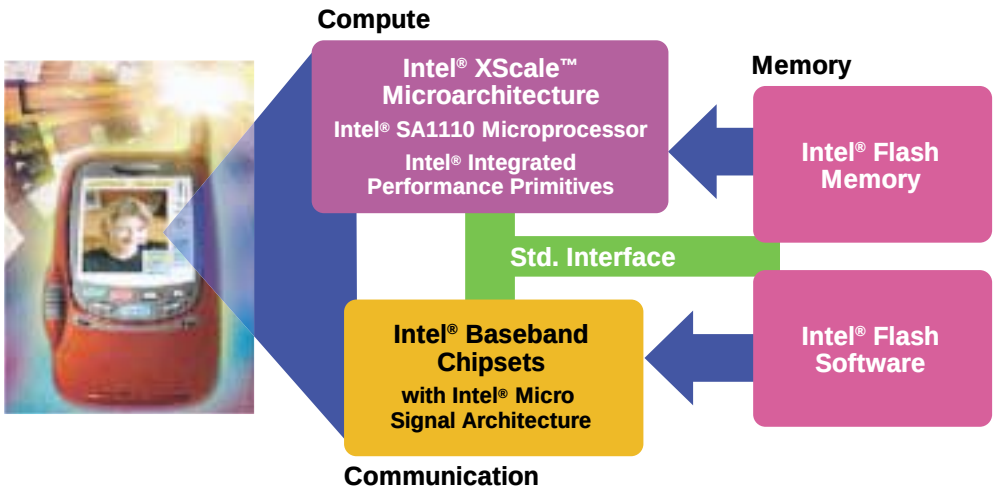

Figure 1. The three main building blocks of the Intel® PCA, linked by a simple standard interface.

Flexibility that Enhances Functionality

As third-generation cellular standards provide the bandwidth to deliver more data, wireless clients are looking to support more exciting Internet content and multimedia applications such as color animation, speech recognition, streaming audio and video, and enhanced remote computing. The Intel PCA helps support speedy development of these demanding applications.

On the hardware side, the Intel PCA enables a high degree of development creativity for voice and data-centric devices. Through its exceptional flexibility, scalability, and modularity, developers can design a phone with more data capability or a PDA with more voice capability, or develop a solution that excels in both capabilities.

The architecture supports a wide range of operating systems and air interfaces, allows integration of custom peripherals, and provides a framework for easily adding and interchanging hardware and software building blocks. These capabilities greatly expand developers' opportunities for creating high-performance Internet products and services, both now and in the future.

Building Blocks that Boost Performance

To introduce exciting new applications, developers must also be able to deliver performance that meets expectations while consuming minimal power. The individual building blocks of the Intel PCA offer leading-edge features in computing, communications, and memory—helping to ensure that new products and services have the speed and power they require to function at their best.

Here's how the key features of each building block address the intensified demands of today's wireless Internet environment:

Compute/Applications Building Block

The computing building block is based on Intel® StrongARM* microprocessor SA1110, as well as the Intel® XScale™ microarchitecture, its successor. It delivers greater processor speed, and its cutting-edge engineering continually balances the requirements of high performance with the need to minimize power consumption on the fly, making it ideal for mobile wireless Internet devices.

Mark Kovandzich, electronics manager for Psion Enterprise Computing, says the compute/applications building block "allows our applications to crank up the CPU speed only when we need the processing power, then to wind it down to minimize power consumption and extend battery life." By delivering this industry-leading MIPS/Watt performance, the Intel PCA enables wireless clients of all kinds to handle solutions, including high-performance multimedia, that have previously been available only on desktops. Figure 2 shows the key features of the computing building block.

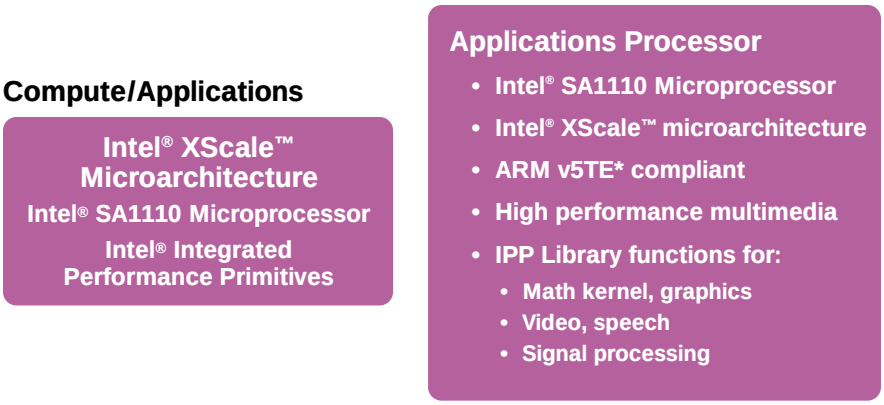

Figure 2. Key features of the computing building block

To support rapid development of applications for Intel StrongARM processors and the Intel XScale microarchitecture, Intel also delivers Intel® Integrated Performance Primitives (IPP) for Intel StrongARM, Intel XScale, and Intel® IA-32 and Itanium™ processors. Intel® IPP enables software vendors, development tools vendors, and original equipment manufacturers to quickly develop exciting new applications that can be ported rapidly to run on any Intel® processor.

Communications Building Block

The communications building block of the Intel PCA is based on Intel® baseband chipsets that incorporate the new Intel® Micro Signal Architecture (Intel® MSA). Intel baseband chipsets are designed into a wide range of cellular phones that support time division multiple access (TDMA) and personal digital cellular (PDC) standards today and will support emerging 2.5G and 3G standards worldwide.

The Intel MSA is a new digital signal processor (DSP) architecture developed by joint development teams from Intel and ADI. It has integrated microcontroller functionality that enables significant improvements in performance, power consumption, and ease of programmability for wireless Internet applications. The architecture also features dynamic power management, which continuously establishes the best blend of performance and power consumption to meet the requirements of real-time applications.

For example, Intel MSA adjusts the power up for video, strikes a balance for voice signal computation, and adjusts performance to maximize battery life in standby mode. As a result, it's ideal for a variety of battery-powered devices that must support high-intensity signal processing with a limited power budget. Figure 3 shows the important features of the communications building block.

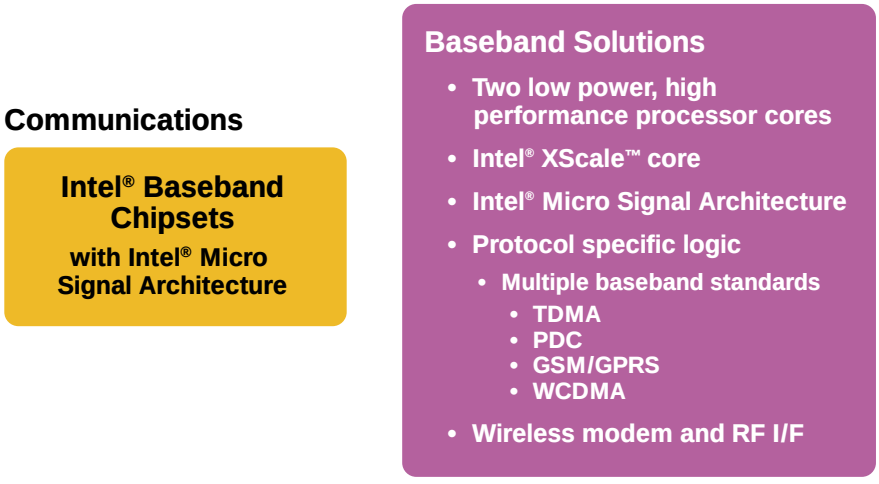

Figure 3. Key features of the communications building block

The communications block also helps reduce software development time because it incorporates a simplified

programming model based on an advanced compiler. With this feature, programmers can code in C/C++. Then, after compilation, areas that have more intense signal processing requirements can be replaced with assembly code through the use of library functions, DSP benchmarks, or handwritten code developed from the signal processor's assembly-level instruction set. In most cases, 80 percent of code will remain in C/C++, substantially reducing development efforts.

Memory Building Block

The memory block of the Intel PCA incorporates Intel® Flash Memory hardware and Intel® Flash Software, so it can take advantage of Intel technology leadership in Flash memory technology to support new levels of performance, with low power consumption.

Intel® StrataFlash™ Memory components utilize reliable two-bit-per-cell technology to deliver high density and lower cost. With the addition of the page mode feature, this technology sets a new benchmark for fast read speeds—up to three times faster than asynchronous reads on previous devices. This feature also provides a high-performance glueless interface to StrongARM SA1110 and future microprocessors.

In addition, Intel® R0 1.8V Wireless Flash memory offers a flexible partition read-while-write operation, with synchronous burst and asynchronous page mode operations. These capabilities enable the performance levels required for voice plus data-enabled wireless devices.

Intel® Flash Data Integrator (FDI) software is a storage manager for use in real-time embedded applications. In wireless clients, FDI and its API can replace EEPROM by handling EEPROM-type data as well as sophisticated functions like streaming data storage.

Intel® Persistent Storage Manager software also helps increase performance while reducing power consumption by combining registry back-up, file storage, and code execution in a single flash component.

Figure 4 shows the important components of the memory building block.

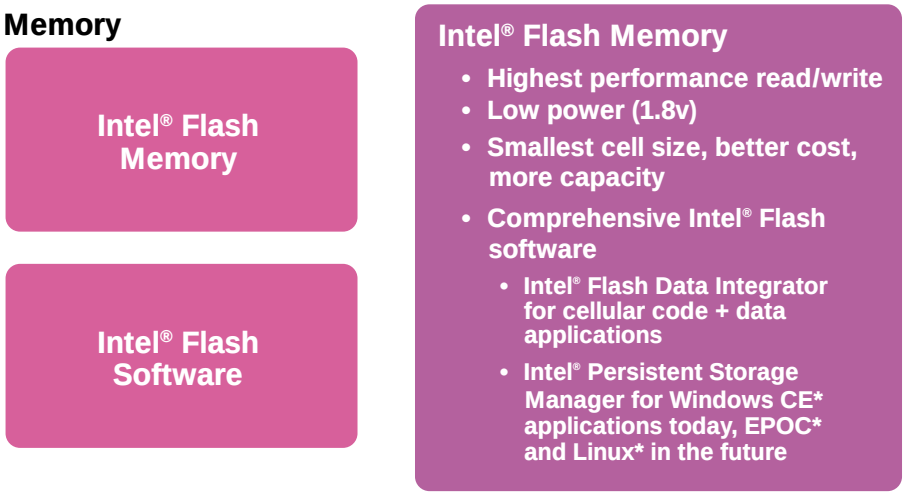

Figure 4. Key features of the memory building block

Benefits to Software Developers

The Intel PCA:

- Eliminates the need to wait for hardware approval before developing software
- Accelerates time-to-market for products
- Allows software applications to be used on a broader range of platforms, from PDAs to smart phones to automotive clients
- Supports creative development of richer software functionality through flexibility, scalability, and modular design
- Helps ensure that product performance meets higher standards required to support new functions

Benefits to Hardware Developers

With Intel Personal Internet Client Architecture and flexible combination of its building blocks, hardware developers can design a single platform that can be used over a broad range of products; cell phones, smart phones, PDAs, and automotive clients.

The Intel PCA:

- Enables better product differentiation through the availability of broad-range applications
- Extends compatibility across most major interfaces to support a global (multiple-market) product offering
- Supports a wide range of processing power and power usage, from very low power platforms to high-end, high-drain performance
- Supports creative development of richer functionality through flexibility, scalability, and modular design
- Helps ensure that product performance meets higher standards required to support new functions

Summary

The Intel Personal Internet Client Architecture (Intel PCA) provides an optimal foundation for developing next-generation wireless Internet devices and applications. This architecture has been specially designed to meet the challenges of today's converging communications and computing environments, and it defines three major building blocks that incorporate all key hardware and software elements for building a new wireless client solution.

By providing a high degree of flexibility, scalability, and modularity, the Intel PCA gives hardware and Internet software developers the capabilities they need to create products, solutions, and services that take full advantage of emerging wireless infrastructure, delivering new levels of performance and functionality. The architecture also provides greater freedom for developers to implement their designs and solutions. For example, with the Intel PCA wireless architecture, hardware developers can design a single platform that can be used over a broad range of products, such as cell phones, smart phones, PDAs, and automotive clients.

The Intel PCA can play a key role in accelerating Internet expansion to the wireless environment—and help hardware and software developers capitalize on today's new product opportunities. Intel is committed to supporting broad use of this architecture across the wireless industry.

Following the launch of the Intel PCA in September 2000, Intel has published a white paper on this architecture and is working with software groups to build an array of development tools that will speed and simplify this implementation process. Developers are encouraged to provide ideas and suggestions for further advances in the Intel PCA.

More Info

For more information about this architecture, refer to the Intel PCA page on the Intel Developer's Web site. The site includes a white paper on the architecture.

The Intel Developer's site also features pages for the XScale microarchitecture, Micro Signal Architecture, and Flash memory components.

Author Bio

Vishwas Deshmane has more than 16 years marketing experience with Intel Corporation, where he has held a wide variety of positions in hardware and software marketing. His product focus has included semiconductor products, and chipsets for PCs, PC peripherals, modems, and LANs. He has also worked as director, communications for the Corporate Marketing Group and as marketing director for the Content Group. He holds an M.S. in physics from University of Bombay and an M.S. in systems management, and an M.B.A. from Florida Institute of Technology.

Making Wireless Networking Easier For Consumers

Adrian M. Ariyanayagam
Communications Networking Initiative Manager
Intel Architecture Group
Intel Corporation

Susan Michalak
Senior Human Factors Engineer
Center For Usability Engineering
Intel Corporation

Overview

Much has been promised by wireless networking and in terms of performance and functionality, much has been delivered. But performance and functionality don't tell the whole story. For the wireless home market, delivering a product that consumers can use easily and confidently is equally if not more essential to success. Networking technology is inherently complex, and wireless networking is no exception; in fact, it adds a layer of complexity to the already mysterious world of home networks. Designers and manufacturers must ensure this complexity does not become an insurmountable barrier to users. This means making usability a top priority.

For the wireless home market, usability challenges can be addressed at three separate points in the product use cycle: pre-purchase and purchase, setup, and ongoing use. At pre-purchase, this means educating distributors, retailers, and users about the unique capabilities of wireless networking. Manufacturers must not assume that removing the wires reduces the complexity for end users. Instead, wireless networking has much of the same complexity as wired networking, and it adds its own set of challenges, especially in terms of interference, interoperability, and security.

Manufacturers also must work to address the end user's perception that wireless is wireless is wireless. So far, little effort has been made to educate distributors, retailers, and users that wireless networking comes in diverse forms that may or may not work in a given environment or with one another. Manufacturers must address these issues clearly and deliberately, in advertising, point-of-purchase collateral, packaging, and end-user documentation.

Mystery of the Missing Information

Internal and external studies alike suggest the most common point-of-purchase problem for users is missing information. A Nortel Networks survey of consumers in March 2000, which can be found on the Ease of Use Roundtable Web site suggests that users cannot easily find the information they need to answer their questions about: (1) what they need to create a complete home networking solution, (2) where various components fit into that solution, (3) what they should expect of plug-and-play components, and (4) what network applications are appropriate for their usage scenario.

To overcome the problem of missing information, past experience dictates that we, as product suppliers, should use multiple mechanisms such as advertising, retail displays, and especially packaging, to inform the consumer of purchase decision factors. Packaging should explain the included components, what additional components will be needed for a complete solution, and how the packaged and additional components will fit into the solution. The same information can and should, where appropriate, be repeated in the user documentation. In addition, as an industry, we should be encouraging the technology press to do articles that discuss compatibility issues such as standards and protocols, and selecting and assembling solutions.

Making Setup Easier

For a wireless home networking environment, both hardware and software setup pose major challenges for users. This is largely because setup involves numerous and often interdependent steps, each of which may involve its own installation and reboot.

Consider the user who may have only a passing knowledge of PC technology, much less wireless networking technology, attempting to set up a home network. There are network cards, drivers, and software components to install on multiple PCs, with a reboot often required after the installation of each component. The user must configure TCP/IP, assign a name to each of the computers in the network, and arrange for file and printer sharing. Finally, the user must determine whether or not everything is working correctly. Through all this, the user must rely on documentation and software installers that often fail to clearly indicate exactly when the process is complete.

It's no wonder that technical professionals have reported that setting up a home network—wired or wireless—can easily take two hours or longer. Imagine the time and patience required of the non-technical user.

To address setup problems, manufacturers must bring more automation to the setup process, such as consistently utilizing "smart" installers to be aware of which components are already installed and which components still need installation. Requesting the end user to find and install the still-needed software components adds significant difficulty and confusion to this task. Additionally, manufacturers should minimize the number of required reboots between installations and make installations self-contained. That is, upon rebooting, software associated with the component just installed would automatically launch the next installer.

Documentation for wireless home networking products should begin with a high-level model of networking and how the various components relate to and locate in the home environment. Much research indicates that non-technical users have no "mental model" of networking, and therefore cannot comprehend how the many elements fit together. Without any end goal in mind, much detailed procedural instruction becomes noise and adds to the confusion.

Once the documentation has explained what networking is and how all the pieces fit within the end user's house, it needs to include detailed overview and procedural information, a clear starting point, and clear distinctions between the steps of setting up various components, and a confirmation of when setup is complete. As an ideal complement to simple and clear documentation, the installation software should offer to demonstrate that the network is configured properly by connecting a client computer to the Internet or printing a page to a shared printer. Documentation needs to be consistent throughout, utilizing detailed and accurate graphics and following through with terminology that is memorable and familiar to target users.

Terminology is a subject worth noting in particular. Wireless networking, just as any other new technology, is generating its share of new terms. A mistaken assumption is that new technology requires new terminology; however, new terminology creates a steep learning curve, especially for home users. To minimize the problem, designers must seek existing or familiar words whenever possible, rather than coining new ones. Another terminology pitfall is making assumptions about what users already know. Testing terminology with end users early on in product development will help avoid these problems.

The Challenges of Security and Troubleshooting

Two areas stand out in terms of the challenges users face during the ongoing use of home networking products: unfamiliarity with relevant security concepts and the lack of a coherent troubleshooting environment. Security, not surprisingly, is a special challenge in the world of wireless technology. A wireless network, with its greater vulnerability to signal interception, needs more elaborate security precautions.

To keep these precautions from overwhelming the home user, manufacturers need to clearly explain the tradeoffs in using various security levels. Describe each of the settings or levels, their benefits and any effects on the use of devices, PCs, or the Internet. For example, explain that if the user sets a higher security setting he or she can expect more obstacles in connecting to the network. Security issues in particular must be documented in a way that is easy for the user to understand. Wherever possible, for example, use of analogies to well-understood concepts (passwords, PINs, modem "handshaking," etc.) can be helpful. Complex security functions (that only an advanced user may care about) should be separated from simple security functions (that a home user may care about).

As for troubleshooting, it is more difficult if users don't have a clear idea of how something works. If a user has a problem but doesn't understand how the product works, the user doesn't know where to begin looking for the cause. Here again, giving the user a high-level "model" of what networks are and how the components communicate with each other will help significantly with troubleshooting later on.

Another key to effective troubleshooting is clear status information. LEDs should be labeled and their meanings clearly explained in documentation. Further, error messages should not only identify a problem but also indicate where the cause might be. Finally, diagnostic tools that test for problems and offer clear steps for troubleshooting (for example, first try *a*, then try *b*…) will help users successfully solve their own networking problems.

Summary

An often-neglected key driver of technology adoption is its usability for the end user/consumer. The functional capability of wireless technology is not enough to deliver on the promise of ubiquitous mobility, especially in the home market. Manufacturers must commit to usability. This will require understanding home users' environment, skills, and needs prior to designing networking solutions for the home, and putting greater effort in providing clear and useful information throughout pre-purchase and purchase, setup, and ongoing use.

As part of this commitment, designers and manufacturers must involve users throughout the product development process to "build in usability" from the beginning. The results—greater user acceptance and lower overall cost of that acceptance—will benefit everyone.

More Info

For more information, visit the Web site of the Ease of Use Quality Roundtable. This is a group that the authors of this article and others at Intel have been involved in to improve usability in a variety of PC-related consumer products. Other information is available through the authors themselves: adrian.m.ariyanayagam@intel.com and susan.michalak@intel.com.

Author Biography

Adrian M. Ariyanayagam began his career with Intel in 1998. With an early focus on software engineering, he has held roles as an application developer in Intel's I/O Products Division and Intel Architecture Labs. Adrian's transition to marketing includes experience in Intel Capital and his current role as technology initiatives manager in charge of communications and networking.

Susan Michalak is a senior human factors engineer and manager in the Center for Usability Engineering at Intel. Before joining Intel in 1999, Susan was a senior human interface designer at Apple Computer, where she focused on hardware and hardware/software integration for a variety of products, including Airport*, Apple's wireless LAN solution.

—End of Intel Developer Update Magazine Issue 18—